DEUTSCHE BUNDESBANK EUROSYSTEM

ChaMP

Discussion Paper

Deutsche Bundesbank
No 38/2026

Granular monetary policy surprises

Falko Fecht
(Deutsche Bundesbank and Frankfurt School of Finance and Management)

Björn Imbierowicz
(Deutsche Bundesbank)

Moritz Kellers
(Deutsche Bundesbank and Goethe University Frankfurt)

Discussion Papers represent the authors' personal opinions and do not necessarily reflect the views of the Deutsche Bundesbank or the Eurosystem.

Editorial Board: Daniel Foos
Stephan Jank
Thomas Kick
Martin Kliem
Malte Knüppel
Christoph Memmel
Hannah Paule-Paludkiewicz

This paper contains research conducted within the network "Challenges for Monetary Policy Transmission in a Changing World Network" (ChaMP). It consists of economists from the European Central Bank (ECB) and the national central banks (NCBs) of the European System of Central Banks (ESCB).

ChaMP is coordinated by a team chaired by Philipp Hartmann (ECB), and consisting of Diana Bonfim (Banco de Portugal), Margherita Bottero (Banca d'Italia), Emmanuel Dhyne (Nationale Bank van België/Banque Nationale de Belgique) and Maria T. Valderrama (Oesterreichische Nationalbank), who are supported by Melina Papoutsi and Gonzalo Paz-Pardo (both ECB), 7 central bank advisers and 8 academic consultants.

ChaMP seeks to revisit our knowledge of monetary transmission channels in the euro area in the context of unprecedented shocks, multiple ongoing structural changes and the extension of the monetary policy toolkit over the last decade and a half as well as the recent steep inflation wave and its reversal. More information is provided on its website.

Deutsche Bundesbank, Wilhelm-Epstein-Straße 14, 60431 Frankfurt am Main,
Postfach 10 06 02, 60006 Frankfurt am Main

Tel +49 69 9566-0

Please address all orders in writing to: Deutsche Bundesbank, Press and Public Relations Division, at the above address or via email: www.bundesbank.de/contact

Internet http://www.bundesbank.de

Reproduction permitted only if source is stated.

DOI https://doi.org/10.71734/DP-2026-38
ISBN 978-3-98848-097-2
ISSN 2941-7503

Granular Monetary Policy Surprises*

Falko Fecht Björn Imbierowicz Moritz Kellers

This version: September 2026

Abstract

Monetary policy surprises are conventionally measured as a single number per announcement, common to all banks. We construct bank-specific surprises from banks' executed overnight index swap transactions in euro area regulatory data, which record identity, price, and intraday timing for every trade around the 36 ECB Governing Council meetings between January 2021 and June 2025. Aggregated across banks, the measure closely tracks established euro area surprise series; within announcements, it reveals dispersion across banks of the same order as the time-series variation of the aggregate surprise. This bank-event variation permits identification that a common shock cannot deliver: firm-month fixed effects absorb common policy news, so we can compare lenders of the same firm in the same month. Banks facing larger surprises reprice newly originated loans to small firms significantly more strongly, by 22 to 38 basis points across the interdecile surprise range over the first two complete post-event months, while the response for large firms is close to zero: the law of one price fails within firm and month. The surprises are costly for the banks themselves: banks with larger absolute surprises experience proprietary equity trading losses of about EUR 1 million per 10 basis points around the same meetings. Corporate deposit rates respond quickly to the surprises; a decomposition attributes this response to common policy news, whereas the bank-specific component prices credit to firms. Monetary policy transmission is thus not uniform across banks: which bank intermediates a small firm's credit shapes the cost of that credit after the same policy decision.

JEL Classification: E52, G10, G21
Keywords: Monetary policy, overnight index swaps, bank heterogeneity, same-borrower identification, loan rates, deposit rates, proprietary trading

*The views expressed here are our own and do not necessarily reflect those of the Deutsche Bundesbank or the Eurosystem. We thank Michael Bauer, Diana Bonfim, Itay Goldstein, Refet Gürkaynak, Rainer Haselmann, Luc Laeven, Maik Schmeling, Michael Weber, and conference and seminar participants at the Frankfurt Summer School 2025, the EMIR Eurosystem Workshop, and the ESCB ChaMP Workshop for helpful comments and suggestions. Fecht: Deutsche Bundesbank and Frankfurt School of Finance and Management, falko.fecht@bundesbank.de. Imbierowicz: Deutsche Bundesbank, bjorn.imbierowicz@gmail.com. Kellers: Deutsche Bundesbank and Goethe University Frankfurt, moritz.kellers@bundesbank.de.

1 Introduction

How strongly monetary policy reaches firms depends on the banks that intermediate it, and a long literature documents that this transmission is heterogeneous across banks and borrowers. Yet empirical studies typically measure monetary policy surprises as a single shock common to all financial institutions (Kuttner, 2001; Gürkaynak et al., 2005; Gertler and Karadi, 2015; Altavilla et al., 2019; Jarociński and Karadi, 2020; Bauer and Swanson, 2023). This approach is informative about aggregate transmission, but a single number per announcement cannot record whether the same decision surprised two banks differently, and firm-time fixed effects absorb a shock that is identical across banks. Applications of the Khwaja-Mian approach therefore interact the common shock with predetermined bank characteristics and identify differential transmission along a particular exposure (Khwaja and Mian, 2008; Kashyap and Stein, 2000; Jiménez et al., 2012, 2014). In this paper, we construct bank-specific monetary policy surprises from banks' own executed overnight index swap (OIS) transactions around ECB Governing Council meetings. The resulting bank-event panel preserves aggregate monetary policy news while retaining the cross-bank variation that aggregate measures average away. We use this variation to examine whether different lenders transmit their own surprises differently to the same borrower after the same policy decision.

Our measurement builds on two regulatory transaction datasets. The Money Market Statistical Reporting Regulation (MMSR) provides reliable contract terms and identifies the reporting bank, while the European Market Infrastructure Regulation (EMIR) supplies intraday execution timestamps. We therefore observe both the identity of the transacting bank and the price at which each OIS trade was executed within the event day, two features that are absent from the aggregated and anonymous interdealer quotes used in conventional high-frequency identification. This information is particularly valuable in the euro area OIS market, where trades are negotiated over the counter and banks need not transact at a single exchange price; recent work on aggregate euro area surprises likewise emphasizes the role of individual participants in OIS rate variation (Altavilla et al., 2025). For 36 Governing Council meetings between January 2021 and June 2025, we compare each bank's pre-announcement transactions with a common post-announcement benchmark for the same maturity and trade direction and aggregate these revisions into one surprise per bank and event.

The transaction data recover the established aggregate signal and reveal substantial variation around it. When we aggregate the underlying transactions across banks before extracting the policy factor, our series has a correlation of 0.80 with the Jarociński and Karadi (2020) interest-rate surprise and 0.72 with their pure monetary policy shock. Averaging the completed bank-specific series yields corresponding correlations of 0.73 and 0.69. At the same time, the completed panel contains 366 bank-event observations: one surprise for each bank at each meeting at which it trades, covering 31 banks and 32 of the 36 meetings. The bank-event surprise has a standard deviation of 6.10 basis

points and an average within-event standard deviation of 3.79 basis points. The latter is of the same order of magnitude as the time-series variation of the aggregate factor over our sample. Our measure thus reproduces the monetary policy news used in high-frequency identification and adds economically meaningful dispersion across banks exposed to the same announcement.

Banks can transact at different pre-announcement rates because they observe different signals, use different models of central-bank behavior, and face different balance-sheet, hedging, inventory, and client-flow needs. Disagreement about the path of policy is well documented (Mankiw et al., 2004; Patton and Timmermann, 2010; Coibion and Gorodnichenko, 2012; Andrade et al., 2019; Bordalo et al., 2020), and recent work relates this disagreement to heterogeneous perceptions of the central bank's reaction function (Bauer et al., 2024; Sastry, 2026). Banks also obtain information through lending and payments and position themselves ahead of policy announcements (Romer and Romer, 2000; Lucca and Moench, 2015; Bernile et al., 2016). We therefore interpret the measure as a realized bank-event transaction-price surprise rather than a direct observation of subjective beliefs. It records the signed difference between the price at which a bank transacted before the policy news and the common market price after the event. This difference can reflect expectations as well as exposures, hedging demand, inventory, client flow, and execution (Bahaj et al., 2025). For our analysis, the key features are that the measure is based on executed positions, reproduces established aggregate policy surprises, and retains the bank-relative variation required for same-borrower identification.

The bank-event panel permits an identification that is unavailable with a common policy shock. We write the observed surprise as the sum of a common event component and a bank-relative deviation, $S_{bt} = C_t + G_{bt}$. Firm-month fixed effects absorb C_t, its interactions with firm characteristics, and all borrower conditions shared by a firm's lenders in a month, including credit demand. The remaining variation, G_{bt}, distinguishes one lender of the same firm from another lender in the same month. Under these fixed effects, the coefficient on the full surprise is algebraically, and in our data numerically, identical to the coefficient on its bank-relative component. A conventional aggregate surprise would instead be absorbed completely. Bank-specific surprises therefore allow us to hold the borrower fixed without eliminating the monetary policy treatment, which expands the set of questions that can be addressed with high-frequency identification. In effect, we ask whether, of two banks lending to the same firm in the same month, the bank that experienced the larger surprise charges more for new credit.

The loan estimates provide our main result. We match the bank-event surprises to newly originated fixed-rate corporate loans in AnaCredit and estimate local projections with bank-firm and firm-month fixed effects. While the pooled specification yields no robust positive average within-firm pass-through, the incidence across firms differs sharply by size. Over the event month and the first two post-event horizons, the response for small firms exceeds the response for large firms

by 3.02, 2.86, and 4.84, respectively, and a synchronized test rejects equality over the five-horizon family. Moving from the 10th to the 90th percentile of the within-event surprise distribution implies 22 to 38 basis points more repricing for small firms in the first and second complete months after the decision. By contrast, the large-firm response remains close to zero throughout.

This result identifies lender-side incidence rather than differences across borrowers. A large literature, beginning with the financial-accelerator evidence of Gertler and Gilchrist (1994), shows that small and bank-dependent firms respond more strongly to monetary policy and emphasizes their weaker access to alternative finance (Kashyap and Stein, 2000; Jiménez et al., 2012, 2014; Crouzet and Mehrotra, 2020; Ottonello and Winberry, 2020; Caglio et al., 2021; Cloyne et al., 2023; Jeenas, 2023). Outside options help explain why firm size matters, but they cannot by themselves generate our estimates because firm-month fixed effects hold the borrower, its credit demand, and all common monthly conditions fixed. Every firm contributing to the comparison borrows from at least two banks in the same month, and identification comes from differences in the surprises of these lenders. The small firms in our design are thus small and medium-sized enterprises that are large enough to maintain several bank relationships, not micro firms; about half of the loan observations involve such multi-bank borrowers. If anything, this restriction works against us: single-bank firms, which cannot enter the comparison, should depend even more strongly on their lender's conditions. We find that different banks charge different rates to the same small borrower after the same policy decision, whereas bank-specific surprises have little effect on the rates offered to large firms. This pattern is consistent with relationship-lending models in which informational advantages and switching costs give incumbent banks more pricing latitude for small firms, while competition and access to market finance limit that latitude for large firms (Sharpe, 1990; Rajan, 1992; Petersen and Rajan, 1995; Degryse and Ongena, 2005; Ioannidou and Ongena, 2010).

Several results support this interpretation. The difference is absent in the remaining days of the event month, emerges in the first full month after the decision, and accumulates through the second month. It is virtually unchanged after controlling for loan amount, instrument type, and maturity. A leave-own-bank-out (LOBO) measure, which excludes a bank's own post-announcement trades from its benchmark, has a level correlation of 0.963 with the original measure and reproduces the same early Small-minus-Large pattern. Positive surprises raise small-firm rates relatively more, while negative surprises lower them relatively more; this direction-consistent pattern, with no detectable asymmetry, speaks against pure rent extraction and for the pass-through of each lender's own assessment into small-firm loan rates. Because the strict specification is identified from 10 or 11 bank clusters, our main inference uses restricted wild-cluster-bootstrap procedures and synchronized horizon-family tests; additional omission exercises are reported in the Appendix.

The surprises also carry economic weight for the banks themselves. Using securities-transaction reports under the Markets in Financial Instruments Regulation (MiFIR), we compute each bank's

proprietary equity trading profit over the seven days up to and including the meeting day. Banks with larger absolute surprises earn significantly lower trading profits around the same meeting: a 10-basis-point absolute surprise is associated with a loss of approximately EUR 1 million. An exact decomposition shows that this association is carried by the common intensity of the event among OIS-active banks. The result validates the measure with a stake computed from banks' own transaction reports, their trading profit and loss, and it complements the loan evidence: common event intensity moves banks' trading books, while the bank-specific component prices credit to firms.

We next examine whether the surprise measure affects a second bank-set price. Active overnight corporate deposit rates respond positively and persistently: a 10-basis-point surprise is associated with about 11 to 12 basis points of cumulative repricing after five to ten business days. In a joint decomposition, deposit rates respond detectably to common event news, while the bank-relative coefficient is positive but imprecise. This contrast is economically informative. Overnight deposit rates in this market are priced off the common overnight benchmark and track common policy news, consistent with the deposits-channel literature (Drechsler et al., 2017, 2021), whereas bank-specific surprises are reflected in the pricing of credit to firms. Finally, we compare Governing Council decisions with narrowly measured policy speeches, using the euro area communication events assembled by Altavilla et al. (2025). When both event types enter jointly, the Governing Council surprises retain the borrower-size incidence, while speech-window surprises do not reproduce it. The granular loan result is therefore concentrated around scheduled policy decisions in our sample.

Our paper contributes to three strands of the literature. First, we extend the measurement of monetary policy shocks. High-frequency identification compresses an announcement into an aggregate time-series innovation (Kuttner, 2001; Gürkaynak et al., 2005; Nakamura and Steinsson, 2018; Altavilla et al., 2019; Jarociński and Karadi, 2020; Swanson, 2021; Bauer and Swanson, 2023), and narrative and information-robust alternatives likewise deliver one shock per meeting (Romer and Romer, 2004; Miranda-Agrippino and Ricco, 2021). We instead construct a bank-event panel from executed OIS transactions while preserving the aggregate signal. This approach differs from using participant quotes to refine an aggregate shock or expand the event set, as in Altavilla et al. (2025), because the participant-level variation itself becomes the treatment in our credit regressions. Our findings also connect transaction-based work on expectations and trading frictions (Bahaj et al., 2025) with evidence on disagreement about monetary policy (Bauer et al., 2024; Sastry, 2026). Regulatory transaction data reveal the heterogeneous bank-level pricing revisions that a single aggregate surprise necessarily averages. The proprietary-trading evidence anchors this measurement contribution in banks' own profit and loss: the surprises we measure are costly for the banks that experience them.

Second, we show that bank-specific surprises permit same-borrower identification without

absorbing the monetary policy shock. The Khwaja-Mian literature uses multiple lending relationships to absorb borrower demand and recover bank credit supply (Khwaja and Mian, 2008; Jiménez et al., 2012, 2014; Amiti and Weinstein, 2018; Degryse et al., 2019). Monetary policy applications commonly interact a single shock with predetermined bank exposures and identify differential transmission along the selected exposure. That approach compares different firms or different bank types under one common shock. Interacting an aggregate surprise with firm size, for example, asks whether small firms respond more than large firms; it cannot ask whether two banks respond differently to the same firm, because a common shock and all its firm-level interactions are absorbed once the borrower-month is held fixed. A saturated triple interaction of the common shock with a bank characteristic and firm size remains feasible within firm-month, as in Jiménez et al. (2014), but the level effect of the common shock is absorbed, so such a design recovers differentials along the chosen characteristic rather than the total response of each borrower group. Because our surprise varies across the lenders of the same firm, firm-month fixed effects can be included without eliminating the treatment and without choosing a particular balance-sheet characteristic in advance, and the total response of each firm-size group is identified. Aggregate shocks measure how the market responds to an event; granular shocks additionally identify how banks exposed to the same event price credit differently to the same borrower.

Third, we contribute to the literature on the incidence of monetary transmission across firms and lending relationships. Granular loan-level evidence shows that bank, firm, and contract components all shape financing conditions and external finance premia (Altavilla et al., 2024; Vilerts et al., 2025), while experimental evidence shows that loan officers with different expectations make different decisions for the same borrower (D'Acunto et al., 2025). We combine market-based policy-event surprises with realized bank loan rates from the euro area credit register. The size differential we estimate is therefore not another firm-characteristic interaction with a common shock; it is identified from differences across the lenders of one and the same firm. Our estimates show that stronger transmission to small firms is not solely a consequence of comparing firms with different credit demand or access to lenders. It remains visible within the same firm and month, when banks with different surprises price credit to the same borrower.

These findings are directly relevant for monetary policy transmission in the euro area, where small and medium-sized firms account for the vast majority of firms and rely predominantly on bank credit. Our estimates imply that the borrowing cost generated by the same policy decision depends partly on the bank that intermediates credit to a small firm, a lender-specific margin that is not detectably present for large firms and that remains invisible with a common surprise measure, even in loan-level data. The measurement results provide a related lesson. Euro area OIS trading is concentrated among a limited number of large institutions, so aggregate surprise series summarize pricing revisions that can differ substantially across banks. Regulatory transaction data make these

granular origins of aggregate fluctuations directly observable (Gabaix, 2011; Gabaix and Koijen, 2024).

The paper proceeds as follows. Section 2 describes the institutional setting and data. Section 3 constructs the surprise measure and connects it to established aggregate series. Section 4 presents the component identity and empirical design. Section 5 reports the loan-rate results. Section 6 shows that the surprises are mirrored in banks' proprietary equity trading. Section 7 examines active deposit repricing and its decomposition into common and bank-relative components. Section 8 compares Governing Council decisions with speeches. Section 9 concludes.

2 Institutional Setting and Data

2.1 OIS Trading around Governing Council Decisions

An overnight index swap exchanges a fixed rate for the compounded overnight reference rate, the euro short-term rate (EUR STR) since 2022 and EONIA before.[1] Because overnight rates are anchored by monetary policy, banks use OIS contracts to manage interest-rate exposure and to take positions on the future path of policy rates. The contracts exchange interest payments rather than principal and are therefore derivatives rather than funding instruments. Euro area OIS trading is closely connected to ECB decisions and has grown rapidly since the exit from the low-rate environment. According to the Eurosystem's money-market statistics, daily turnover reached approximately EUR 228 billion after the OIS segment had roughly doubled between 2022 and 2024. Most trades are forwards, about 60 percent of forward activity is aligned with reserve-maintenance periods and Governing Council dates, and trading activity on meeting days is about twice the daily average (ECB, 2025).

The OIS market is organized over the counter. Trades are commonly initiated on request-for-quote platforms or through voice brokers and subsequently novated to a central counterparty. Consequently, there is no single exchange price at which all banks transact. Execution times, counterparties, balance-sheet exposures, inventories, hedging needs, and assessments of policy can produce different transaction prices even for similar contracts. Recent work that refines aggregate euro area surprises also points to the potentially large role of individual participants in OIS rate variation (Altavilla et al., 2025). We exploit the bilateral and executed nature of these transactions to construct bank-specific monetary policy surprises.

Several economic forces can generate cross-bank differences in executed prices. Banks obtain different real-time signals about economic activity through their lending relationships and payment flows (Romer and Romer, 2000); they may use different models of the central bank's reaction

[1] The transition from EONIA to EUR STR was officially completed on January 3, 2022, though the majority of OIS trades had already shifted to EUR STR referencing by mid-2021.

function (Mankiw et al., 2004; Andrade et al., 2019; Bauer et al., 2024; Sastry, 2026); and their balance-sheet exposures shape hedging and positioning in the OIS market (Kashyap and Stein, 2000; Begenau et al., 2015; Hoffmann et al., 2019; Gomez et al., 2021). Trade itself presupposes that counterparties differ in information, hedging needs, or beliefs (Milgrom and Stokey, 1982). Accordingly, we do not interpret all cross-bank price dispersion as disagreement about monetary policy. Our measure captures the change from each bank's pre-announcement transaction price to a common post-announcement benchmark and records the bank's realized transaction-price surprise around the policy event.

Our event sample contains 36 ECB Governing Council meetings from January 2021 through June 2025. It covers the end of the effective-lower-bound period (Heider et al., 2019), the 2022–2023 tightening cycle, the subsequent plateau, and the beginning of easing. Until June 2022, policy decisions were released at 1:45 p.m. Central European Time; since July 2022, the release has occurred at 2:15 p.m., with the press conference and question-and-answer session ending around 3:30 p.m. and 3:45 p.m., respectively. We classify transactions before the release as pre-event trades and transactions after the end of the press conference as post-event trades. We use the full pre-event trading day because executed trades are discrete events rather than continuous quotes: activity peaks in the late morning and becomes thin immediately before the release. We similarly allow post-event trading to resume after the press conference, where activity peaks about one hour later, rather than imposing a narrow quote-based window. Appendix Figure A.1 documents these intraday patterns.

The sample period contains exceptional monetary policy variation. Between July 2022 and September 2023, the ECB raised its deposit facility rate from −0.50 to 4.00 percent, the fastest tightening in the history of the euro, before the first rate cuts in June 2024. Over the sample period, excess liquidity fell from a peak of roughly EUR 4.7 trillion in late 2022 to below EUR 3 trillion in late 2024 and about EUR 2.7 trillion by mid-2025 as targeted longer-term refinancing operations matured. This variation is valuable for our purposes: it generates policy surprises of both signs and substantial magnitude, and it implies that the measured bank-event dispersion does not reflect a single quiet policy regime.

2.2 Regulatory Transaction Data

We combine two reporting frameworks. MMSR is the ECB's core money-market dataset: the largest euro area reporting banks submit their daily money-market transactions, including OIS trades, with contract rates, notionals, maturities, trade direction, and counterparty information. MMSR is our source of record for the economic terms of each contract, but it does not contain intraday execution timestamps. EMIR requires euro area counterparties to report all over-the-counter derivative transactions to trade repositories. It covers a broader universe of market participants

and contains execution timestamps, although the reports include more errors and duplicates. We link both datasets through unique trade identifiers, reconcile counterparties and directions, remove duplicates and novations, and retain centrally cleared OIS contracts with maturities from one week to ten years. Appendix A.1 describes the cleaning sequence in detail. The matched data identify the transacting banks and provide executed prices with the intraday timestamps required to distinguish pre- and post-announcement transactions.

The cleaned event-day panel contains 8,171 OIS transactions: 5,210 before the policy event, 1,235 during the release and press conference, and 1,726 after it. The linked sample covers the concentrated core of the euro area OIS market. The ten most active MMSR banks account for 95 percent of aggregate turnover (ECB, 2025); in EMIR, approximately 90 percent of reports involve at least one bank and 84 percent involve at least one MMSR reporting bank as counterparty. The completed surprise panel contains 366 bank-events from 31 banks; it spans 32 of the 36 meetings, since at the remaining four the transaction support is insufficient to construct a bank-event surprise. The strict loan regressions are identified from 10 or 11 bank clusters because they additionally require same-month multiple lending relationships, lagged outcomes, and fixed-effect support. We report the effective number of banks in every main table, and Appendix Table A.1 traces the sample from the reporting universe to each estimation sample.

2.3 Corporate Loans and Deposits

Loans. We measure loan pricing using AnaCredit, the euro area credit registry. We begin with fixed-rate term loans and non-revolving credit lines extended by banks in the OIS sample to nonfinancial corporations between January 2021 and June 2025. We retain euro-denominated, newly originated instruments with positive rates and residual maturity of at most one year. For credit lines, commitment amounts replace nominal amounts where appropriate. We exclude defaulted debtors, reconcile parent entities, remove repeated reports, and aggregate transaction rates to the bank-firm-month level using loan amounts as weights. These restrictions yield a relatively homogeneous set of contracts. Newly agreed fixed rates reflect bank pricing decisions rather than mechanical resets, and the contract horizon overlaps the short and medium OIS maturities used in the surprise measure. The restrictions also reduce pass-through heterogeneity arising from rate-fixation details and reference-rate choice, which euro area evidence shows to be quantitatively important (de Bondt, 2005; Altavilla et al., 2020; Vilerts et al., 2025). The unrestricted cleaned loan source contains 943,288 observations from 27 banks and 107,170 firms, and roughly half of the observations involve firms that borrow from several banks. The same-borrower samples are smaller because the regressions require multiple lenders, lagged outcomes, and fixed-effect support. To embed the bank-event measure in the monthly loan panel, we follow the convention of

a monthly shock panel and assign a zero surprise to bank-months without a measured bank-event surprise, including months without a meeting and meetings at which a bank reports no qualifying OIS transaction. This is a zero-exposure convention: a bank that does not transact in the OIS market around a meeting reveals no pricing revision there, and for the banks in the equity sample nine in ten such bank-events involve no qualifying pre-announcement trade at all (Appendix A.4).

We classify firm size using lagged annual turnover reported in AnaCredit. A firm is large when its lagged turnover is above the sample median and small otherwise. The indicator is lagged to prevent contemporaneous loan pricing from determining group assignment. Our estimand compares the shock slopes for these two mutually exclusive groups in the same joint regression. Firm size is a transparent proxy for the availability of outside finance, not a label for financial weakness; size can reflect many dimensions other than constraints (Crouzet and Mehrotra, 2020). Appendix evidence based on continuous lagged turnover is directionally consistent but less precise.

Deposits. MMSR also records unsecured money-market borrowing by reporting banks. We use newly reported fixed-rate overnight deposits for which a sample bank is the borrower and the ultimate counterparty is a nonfinancial corporation. We require same-day settlement, positive amounts, no notice period or embedded option, and consistent counterparty information, and we consolidate counterparties at the ultimate parent (Fecht et al., 2025). Multiple transactions within a bank-firm-day are volume weighted. The main outcome focuses on active relationships whose rate or volume changes relative to the previous business day, because unchanged carried-forward reports do not represent a new pricing decision. The source panel contains 631,670 observations from 30 banks and 18,886 nonfinancial firms; overnight is by far the most common maturity, and about half of the observations involve depositors with several bank relationships. The main deposit regressions include bank-depositor and depositor-by-Governing-Council-cycle fixed effects.

The two samples address complementary questions. Monthly loan originations permit a strict same-borrower test of lender-side incidence, while daily deposit transactions show how the full surprise measure and its two components relate to a denser bank-pricing margin. MMSR deposit counterparties cannot be matched to AnaCredit borrowers at the firm level, so the two designs are deliberately self-contained; none of our conclusions requires such a link.

3 Measuring Granular Monetary Policy Surprises

3.1 Cross-Bank Dispersion in Executed Prices

The raw data already display substantial cross-bank variation. Within the same maturity bucket and Governing Council day, banks trade at different fixed rates before the policy announcement. The

average cross-bank standard deviation of volume-weighted pre-announcement rates is 4 basis points in the short-term bucket (up to three months), 9 basis points in the medium-term bucket (three months to one year), and 17 basis points in the long-term bucket (one to ten years).[2] Appendix Figure A.2 shows that the amount of dispersion varies across maturities and phases of the policy cycle. We next translate the observed differences in price levels into bank-specific counterparts of a conventional event-study surprise.

3.2 Transaction and Bank-Event Construction

We construct monetary policy surprises in two steps, following standard high-frequency identification but applying it to executed transactions at the bank level (Kuttner, 2001; Gürkaynak et al., 2005; Altavilla et al., 2019; Jarociński and Karadi, 2020). For pre-announcement transaction i by bank b on event date t, narrow maturity band m (one week through ten years), and trade direction d (pay or receive fixed), define

$$s_{ibtmd} = r^{\text{post}}_{tmd} - r^{\text{pre}}_{ibtmd}. \tag{1}$$

The first term is the volume-weighted post-announcement fixed OIS rate across all banks in the same event, maturity band, and pay- or receive-fixed direction. The second is the executed rate on bank b's pre-announcement trade. Positive values correspond to an upward revision in the OIS rate. Matching on direction prevents bid-ask differences from mechanically becoming surprises; narrow maturity bands prevent term-structure differences from doing so; and volume weighting the benchmark limits the influence of very small tickets.

Every pre-event transaction is measured against a common post-event benchmark. A bank's own post-event prices may reflect adjustments after the announcement, including unwinding, hedging, inventory management, and client flow. Using these prices as bank-specific benchmarks would therefore combine the initial surprise with the bank's subsequent response. The pooled benchmark instead evaluates all pre-event transactions against the same post-event price for a comparable contract, so that cross-bank variation originates on the pre-announcement side. The economic rationale is that the complete communication of the decision, including the press conference, resolves the uncertainty that generated pre-announcement disagreement: cross-bank rate dispersion falls markedly once the press conference has ended (Appendix Figure A.2), so the volume-weighted market price is the natural common reference for what the decision implied. As a robustness test,

[2] We first calculate volume-weighted averages of the fixed rate within bank, maturity bucket, and Governing Council day, then compute the standard deviation across banks for each bucket-day and average over days. When we additionally condition on trade direction, grouping pay-fixed and receive-fixed contracts separately in both steps, the resulting standard deviations are very similar for the two directions and virtually unchanged in size, indicating that cross-bank variation is not mainly driven by bid-ask spreads. Even within very narrow maturity bands such as two months, two years, or ten years, the cross-bank standard deviations average two, seven, and six basis points, respectively, for both trade directions.

the LOBO construction in Appendix Table A.2 excludes bank b from its own post-event benchmark. Both measures are nearly identical on common support, showing that a bank's mechanical influence on the pooled benchmark is quantitatively unimportant.

We next volume weight transaction-specific surprises within bank, event, and three broad maturity buckets: up to three months, three months to one year, and one to ten years. In the spirit of the interest-rate surprise of Jarociński and Karadi (2020) and Jarociński (2022), we extract the first principal component from the short and medium buckets and normalize it to the medium bucket for interpretability. The resulting bank-event measure, denoted S_{bt}, summarizes the revision in the segment of the OIS curve most closely connected to the target and near-term path of policy, the maturities associated with the target and timing factors in the literature on the dimensions of monetary policy (Gürkaynak et al., 2005; Altavilla et al., 2019). A bank-event surprise enters the measurement panel only when the underlying trades are observed; the 366 bank-events are measured, not imputed. Section 2 describes how the measure is embedded in the monthly loan panel.

An individual transaction need not reveal a bank's expectation exactly because its price also reflects counterparty matching, hedging, inventory, and execution. Averaging the approximately six pre-event prices observed per bank and event reduces transaction-specific variation. We interpret S_{bt} as a realized bank-event transaction-price surprise: the signed change from the price at which a bank transacted before the policy news to a common price after the event. Its aggregate component can be compared with established monetary policy surprises, while its bank-relative component provides the cross-bank variation used in the loan regressions.

3.3 Aggregate Continuity and Granular Support

Figure 1 provides the first validation of the measure. We aggregate transactions across banks before extracting the principal component and compare the resulting series with Jarociński and Karadi (2020), whose factors are constructed from indicative interdealer quotes (Altavilla et al., 2019); OIS rates at these maturities primarily reflect expectations about the policy path rather than risk premia (Schmeling et al., 2022). The correlation is 0.80 with their interest-rate surprise and 0.72 with their sign-restricted pure monetary policy shock.[3] These correlations are remarkably high given that our measure is based on executed prices in regulatory reports for 36 event days, whereas the benchmark uses continuous indicative quotes from a commercial data vendor. When we first construct the bank-event surprises and then average across banks, the corresponding correlations are 0.73 and 0.69. The first comparison is closest to conventional aggregate high-frequency identification; the second confirms that the aggregate policy signal remains present after retaining bank identity. Our objective is therefore not to replace existing aggregate series, but to recover the cross-bank variation

[3]The overlapping period runs from March 2021 to April 2025. The correlations are particularly strong through the end of 2023, during the active tightening phase, at approximately 0.90 and 0.78, respectively.

underlying them.

Table 1 reports the same bridge numerically and adds the transaction-level distribution. The 4,127 transaction-specific surprises have a standard deviation of 11.06 basis points. Aggregation reduces idiosyncratic transaction variation, but the completed bank-event measure still has a standard deviation of 6.10 basis points and spans roughly 55 basis points. The average within-event standard deviation of 3.79 basis points is large relative to the approximately 5-basis-point time-series standard deviation of the Jarociński-Karadi factor over the overlapping sample. The data therefore contain two dimensions: common variation over events and bank-relative variation within events. The mean bank-event surprise is close to zero, consistent with bank-relative components largely averaging out across banks, which is also why the aggregated series tracks the established benchmarks.

Two additional diagnostics examine whether the cross-bank variation reflects observable bank types or the construction of the benchmark. First, Appendix Table A.3 relates signed and absolute surprises to predetermined bank size, capitalization, wholesale deposit dependence, and nonfinancial-corporate deposit dependence. We distinguish the signed surprise, S_{bt} itself, from its absolute value: bank characteristics could predict the direction of a bank's surprise, or they could predict its typical magnitude, for instance if particular business models produce larger pricing errors in either direction, and we test both. With event fixed effects, the joint WCR-S p-values are 0.268 for signed surprises and 0.265 for absolute surprises.[4] Individually, only two associations appear: wholesale deposit dependence is marginally related to the signed surprise, and larger banks display smaller absolute surprises, by 0.44 basis points per standard deviation of log assets, consistent with the averaging of more transactions per bank-event; neither survives as a joint pattern, and the loan identification in any case relies on signed within-firm-month variation after absorbing bank-firm effects. Adding bank fixed effects yields joint p-values of 0.731 and 0.864, respectively. Thus, stable observable bank types have little joint power for predicting the surprises, either across banks within an event or over time within banks: the surprises are not a relabeling of familiar balance-sheet characteristics. Second, the original and LOBO measures correlate 0.916 within events and have the same within-event sign for 93.4 percent of common bank-events (Appendix Table A.2). We use the original measure throughout the main analysis and report LOBO as a robustness test.

Taken together, the results show that the transaction data preserve established aggregate monetary policy news and add substantial variation among banks exposed to the same event. The next section explains how this cross-bank variation permits same-borrower identification in bank-firm data.

[4]Throughout the paper, WCR-S denotes p-values from the restricted wild cluster bootstrap with Webb weights (MacKinnon et al., 2023), based on 9,999 draws and clustered at the bank level unless stated otherwise. The bootstrap imposes the null hypothesis when resampling and is the recommended inference procedure when the number of clusters is small.

4 Identification and Empirical Strategy

4.1 Common and Bank-Relative Components

Let S_{bt} denote the observed surprise of bank b at Governing Council event t. Each bank's surprise consists of a component that is common to all banks at the event and a bank-specific component:

$$S_{bt} = C_t + G_{bt}, \tag{2}$$

where C_t is the average surprise across the banks observed at event t and G_{bt} is bank b's deviation from that average. The common component corresponds to the object that aggregate high-frequency series measure; the bank-specific component is the additional variation that our data contain. The decomposition is an accounting statement and does not separately identify beliefs, hedging demand, or trading frictions; its purpose is to make the source of identifying variation explicit.

Suppose loan rates are observed for bank b, firm f, and event month t. A conventional aggregate shock has no cross-bank variation and is absorbed by firm-month fixed effects. Monetary policy applications in the Khwaja-Mian tradition therefore interact the common shock with predetermined bank exposures such as liquidity, capitalization, or funding structure (Kashyap and Stein, 2000; Jiménez et al., 2012, 2014). These specifications identify differential transmission along the selected exposure. They can be saturated with firm-month fixed effects when the exposure varies across a firm's lenders, as in the triple interaction of Jiménez et al. (2014); the level effect of the common shock is then absorbed, so the estimates remain differentials along the chosen exposure. Bank-specific surprises provide a different source of variation. Since firm-month fixed effects absorb C_t, the coefficient on S_{bt} is algebraically the coefficient on G_{bt}. We verify this identity at every horizon: replacing the original surprise with the bank-relative component leaves the Small-minus-Large coefficient unchanged to reported precision. Under the zero-exposure convention of Section 2, a lender without a qualifying OIS transaction at a meeting has a surprise of zero and hence a bank-relative deviation of $-C_t$; the same-borrower comparison therefore contrasts lenders that revised their pricing in the OIS market with lenders of the same firm that did not, as well as measured lenders among themselves. The granular measure therefore permits a comparison of lenders to the same firm after absorbing common policy news and borrower-month conditions, without selecting a particular balance-sheet exposure in advance, and it identifies the total response of each borrower group rather than a differential along a chosen exposure.

4.2 Local Projections and the Incidence Estimand

We estimate monthly local projections (Jordà, 2005). For horizon h, the pooled same-borrower specification is

$$R_{bf,t+h} - R_{bf,t-1} = \beta_h S_{bt} + \Gamma_h' X_{bf,t-1} + \rho_{bf}^h + \alpha_{ft}^h + \varepsilon_{bf,t+h}, \tag{3}$$

where $R_{bf,t+h}$ is the amount-weighted rate on newly originated short-maturity fixed-rate loans, ρ_{bf}^h are bank-firm fixed effects, and α_{ft}^h are firm-month fixed effects. $X_{bf,t-1}$ contains two lags of the outcome and surprise terms. The dependent variable is cumulative relative to month $t-1$. The pooled coefficient asks whether lenders with larger bank-relative surprises charge a higher rate to the same borrower on average. Appendix Table A.5 shows that the data do not deliver a robust positive average within-firm response.

Our primary estimand is incidence by firm size. Let $D_{f,t-1}^S$ and $D_{f,t-1}^L$ be mutually exclusive indicators for turnover at or below and above the lagged sample median. We estimate

$$\begin{aligned} R_{bf,t+h} - R_{bf,t-1} = \beta_h^S S_{bt} D_{f,t-1}^S + \beta_h^L S_{bt} D_{f,t-1}^L + \Gamma_h' X_{bf,t-1} \\ + \rho_{bf}^h + \alpha_{ft}^h + \varepsilon_{bf,t+h}. \end{aligned} \tag{4}$$

The two coefficients are total slopes for the two groups and are estimated jointly. Since the surprise does not enter separately, each interaction captures the full response for its firm-size group. Our parameter of interest is the post-estimation contrast $\Delta_h = \beta_h^S - \beta_h^L$, which tests whether loan rates for the same small borrower respond more strongly to differences in lender surprises than loan rates for the same large borrower. Because $C_t D_{f,t-1}^S$ and $C_t D_{f,t-1}^L$ vary only at the firm-month level, both are absorbed. The contrast is therefore identified entirely from bank-relative surprise variation.

For a causal interpretation, we require that, conditional on the fixed effects and controls, bank-relative OIS pricing surprises are uncorrelated with contemporaneous bank-firm-specific shocks to new-loan pricing. Firm-month fixed effects remove all borrower conditions shared by a firm's lenders in a month, and bank-firm fixed effects absorb permanent features of each lending relationship. Time-varying lender-firm-specific demand may nevertheless remain when a firm uses different banks for different purposes (Paravisini et al., 2023). We address this concern by restricting the outcome to newly originated euro-denominated fixed-rate credit with maturity of at most one year and by testing loan amount, instrument type, and maturity directly in Section 5. Moreover, the post-event benchmark is common within maturity band and event day, so any error in that benchmark is absorbed by firm-month fixed effects; identification comes from differences in pre-announcement transaction prices across lenders. Under the identifying assumption, the estimates measure differential lender-side repricing. They do not distinguish whether banks adjust through

funding costs, risk limits, internal transfer prices, inventory, or another balance-sheet margin.

4.3 Economic Predictions

Relationship-lending theory provides a natural interpretation for heterogeneous pass-through across firms. Large firms have broader access to market finance and face competition among a larger set of lenders, which limits the extent to which an individual bank can incorporate its own conditions into the loan rate. For small firms, screening and switching frictions can create informational advantages and greater pricing latitude within an incumbent relationship (Sharpe, 1990; Rajan, 1992; Ioannidou and Ongena, 2010). Consistent with this argument, loan rates vary with local banking competition and the lender's informational position (Petersen and Rajan, 1995; Degryse and Ongena, 2005). The newly originated fixed-rate loans in our sample have maturities of at most one year, so their pricing horizon overlaps the OIS maturities used to construct the surprises. If banks incorporate their own policy-event pricing revisions into new credit, the response should therefore be stronger where lender-specific pricing latitude is greater.

This interpretation yields four empirical implications. First, the Small-minus-Large contrast Δ_h should be positive at early post-event horizons. Second, it should appear most clearly in loans originated after the complete policy event rather than in the residual days of the event month, when previously negotiated loans may still settle. Third, the contrast should be stable after controlling for observed contract composition. Fourth, the response should be direction consistent: positive surprises should raise small-firm rates relatively more, while negative surprises should lower them relatively more. This fourth implication deserves explanation, because it is not what a pure rent-extraction view of lending relationships would predict. If incumbent banks merely exploited informational lock-in to extract rents, rates should rise readily after positive surprises but fall reluctantly after negative ones. If, instead, banks price new credit from their own assessment of funding conditions and small firms are the borrowers whose rates follow their own lender's assessment rather than a market-wide benchmark, the pass-through carries both signs: the relationship determines whose assessment is priced, not the direction of the adjustment. The sign split therefore distinguishes between these two readings of relationship lending. All four tests compare lenders to the same firm in the same month and therefore remove differences across borrowers as an explanation for the estimated size differential. They are designed to assess lender-side repricing, while the specific internal bank mechanism remains outside the scope of the identification.

4.4 Inference

The strict design contains a small number of bank clusters, and with few clusters conventional clustered standard errors reject far too often. The loan and equity tables therefore report restricted wild-cluster-bootstrap (WCR-S) p-values based on 9,999 Webb draws and clustering by bank; the restricted bootstrap imposes the null hypothesis when resampling, which is designed to restore size control in exactly this situation (MacKinnon et al., 2023). The deposit tables state their clustering procedures in the notes. The coefficient figure additionally shows conventional bank-clustered confidence intervals to display magnitudes and sampling dispersion. Because our result is a path across horizons rather than a single coefficient, we complement the pointwise tests with synchronized horizon-family tests, which ask whether the path as a whole differs from zero, and with leave-one-bank and leave-one-event exercises in the Appendix. We evaluate the results jointly on the basis of their sign, economic magnitude, measurement robustness, and few-cluster inference.

5 Granular Surprises and Loan Repricing

5.1 From Conventional Saturation to the Same-Borrower Design

We first estimate the specifications that correspond to the conventional approach and then show what the stricter design identifies. Panel A of Table 2 includes only bank-firm fixed effects. In this setting, the common and bank-specific surprise components can be estimated jointly, and we report both the pooled slopes and their small-firm interactions. The estimates are uninformative: coefficients are small, switch signs across horizons, and none is statistically significant. This outcome has a clear reason. Without firm-month fixed effects, borrower-month demand remains uncontrolled, and the common component must be identified from the time series of only 36 events. The conventional saturation therefore cannot answer how different banks price credit to the same firm.

Panel B adds firm-month fixed effects and delivers our main result. Each column is one joint regression; the first two rows are the total surprise slopes for small and large firms, and the Small-minus-Large difference is reported as a post-estimation contrast. The response for small firms is positive at every horizon and strongest in the first three months, while the response for large firms is close to zero throughout. The Small-minus-Large difference equals 3.018 in the event month, 2.863 in the first full post-event month, and 4.845 through the second full month, with WCR-S p-values of 0.041, 0.049, and 0.034. It subsequently declines and is statistically insignificant at horizons three and four. The comparison across the two panels summarizes the identification argument of Section 4 in one table: once borrower-month demand is absorbed and the bank-specific component is the only remaining treatment variation, a sharp and economically interpretable pattern emerges from estimates that appear scattered under conventional saturation.

Because the result is a path over five horizons rather than a single coefficient, we also test the horizons jointly. A synchronized family test asks whether the five Small-minus-Large differences are jointly zero, which guards against the concern that an individual significant horizon reflects chance under multiple testing. The test rejects with a p-value of 0.031 on the common set of banks observed at all horizons; a more conservative version that fixes the observation support across horizons yields 0.082. The pooled same-borrower average, in contrast, is not robustly positive: small firms respond, large firms do not, and the average mixes the two. Appendix Table A.5 reports the pooled estimates with confidence intervals; at horizon one, a 10-basis-point surprise carries a 95 percent interval from -21.0 to 5.0 basis points.

The estimated differences are economically substantial. Evaluated at the horizon-specific 90th-to-10th-percentile spread of within-event surprises, they imply 23.5 basis points more repricing for small firms in the event month, 22.3 basis points after the first full month, and 37.8 basis points through the second full month. For firms that rely on bank credit as their main source of external finance, these differences in the price of newly originated short-maturity credit are sizable. The calculations compare banks at different positions in the observed cross-bank surprise distribution while holding the borrower-month fixed.

Figure 2 presents the paper's central result graphically. Panel A traces the two total slopes over the five horizons. The small-firm path lies above zero from the event month onward, peaks at the two-month cumulative horizon, and declines thereafter; the large-firm path is flat and statistically indistinguishable from zero at every horizon. The gap between the two paths is the paper's estimand, and the p-values printed above the paths test its significance directly at each horizon; the confidence bands shown are conventional bank-clustered intervals around each slope. Panel B translates the gap into economic magnitudes: it plots, for each horizon, the additional repricing of small firms implied by moving a lender from the 10th to the 90th percentile of the within-event surprise distribution. The figure conveys the paper's message in compact form. Two banks that walked into the same Governing Council meeting with different positions charge the same small firm rates that differ by more than 20 basis points in the following months, while a large firm borrowing from the same two banks faces virtually identical conditions.

The strict specification is identified from 10 or 11 banks because it requires the same borrower to obtain comparable new loans from several lenders in the relevant month. The pooled same-borrower specification contains 28,834 observations from 12 banks at horizon zero; requiring lagged turnover to classify firms as small or large leaves 26,129 observations from 11 banks in the firm-size specification reported in Table 2. Appendix Table A.6 therefore reports the complete inference battery. The early Small-minus-Large coefficient remains positive when any single bank or any single event is removed, the leave-own-bank-out measure reproduces it, and the family tests support the joint path. Sign, economic magnitude, measurement robustness, and few-cluster inference point

in the same direction.

5.2 When Does the Differential Emerge?

The second prediction concerns timing: the differential should appear in credit originated after the complete policy event, not in loans that were already negotiated before the decision. The monthly outcome cannot separate the two directly, because a meeting held partway through a month leaves loans originated both before and after the decision in the same event month. Table 3 therefore splits the event-month window. The strict column compares the rate on loans originated strictly after the Governing Council date, but still within the event month, with the rate in the previous month; the inclusive column adds loans originated on the meeting day itself. The remaining columns report the cumulative changes through the first and the second complete post-event month, as in Table 2.

The Small-minus-Large difference is close to zero in both residual event-month windows and becomes positive and significant in the first full month, accumulating through the second. The horizon-zero estimate in Table 2 should be read in this light: it averages over all loans originated in the event month, before and after the meeting, and is therefore a mixed rather than a strictly post-decision estimate, whereas the strict window isolates the few days between the decision and the month end and retains less than half of the observations. The clean post-event evidence begins in the first complete month after the decision and accumulates through the second. New credit is thus repriced with a short delay rather than within days: rate offers outstanding at the decision date appear to be honored, and the full differential materializes once lending is entirely post-event. The economic magnitudes we emphasize throughout, 22 to 38 basis points over the interdecile surprise range, refer to these two fully post-event horizons.

5.3 Contract Composition

The third prediction is that the differential reflects pricing rather than changes in the composition of credit. Firm-month fixed effects hold the borrower fixed but do not force every newly originated contract to have identical terms: if small firms obtain different maturities, amounts, or instruments after particular events, the estimated incidence could partly reflect changes in the credit bundle. Table 4 addresses this concern. Each row re-estimates the complete horizon path, first without additional contract controls, then controlling for amount, instrument type, and mean residual maturity, and finally replacing linear maturity with maturity-bin indicators.

The estimates are remarkably stable. At horizons zero through two, the richest specification yields coefficients of 2.962, 2.891, and 4.873, compared with baseline values of 3.018, 2.863, and 4.845, and sample attrition is negligible. Differences in observed loan amount, instrument type, and maturity therefore do not account for the early size differential. The result reflects different rates on

observably comparable credit rather than a change in the bundle of contracts originated for small firms.

5.4 Positive and Negative Surprises

The fourth prediction is that the differential operates in both directions: banks with positive surprises should raise small-firm rates relatively more, and banks with negative surprises should lower them relatively more. This prediction distinguishes between two readings of relationship lending. If incumbent banks simply exploited informational lock-in to extract rents, one would expect an asymmetric pattern in which rates rise readily but fall reluctantly. If, instead, banks price new credit from their own assessment of policy and funding conditions, and small firms are the borrowers whose rates follow their own lender's assessment rather than a market-wide benchmark, then the pass-through should carry both signs. In the second reading, the relationship determines whose assessment is priced, not the direction of the adjustment; competition for new loans, which all borrowers in our multi-lender sample can in principle invoke, disciplines the level of the markup but leaves the lender-specific revision visible (Sharpe, 1990; Rajan, 1992; Petersen and Rajan, 1995; Ioannidou and Ongena, 2010).

The sample spans both upward and downward OIS revisions, so we can test this directly. We split the signed surprise into positive and negative parts and include both in the same regression, together with their firm-size interactions. Table 5 reports the Small-minus-Large slope for each part. This is a sign decomposition of the bank-event surprise, not a comparison of historical tightening and easing regimes.

The coefficient is positive for both signs at every reported horizon. A positive surprise is associated with a larger rate increase for small firms, whereas a negative surprise is associated with a larger rate decline. Both slopes are statistically informative at horizons one and two, and direct equality tests yield p-values between 0.164 and 0.195. We therefore find direction-consistent repricing and no statistically detectable asymmetry. The evidence favors the pass-through reading over pure rent extraction: small firms are not merely charged more when their bank is surprised upward, their loan rate tracks their own bank's realized surprise in both directions. This also means that small firms bear the dispersion of lender assessments rather than a one-sided markup, an implication we return to in the conclusion.

5.5 Economic Interpretation

Tables 2–5 show that bank-specific surprises pass through more strongly to small firms than to large firms borrowing in the same month. The difference is concentrated in the early post-event horizons, remains stable after controlling for observed loan characteristics, and has the economically

consistent sign after both positive and negative surprises.

The same-borrower design adds a lender-side pricing margin to the established evidence on small-firm monetary transmission. Since Gertler and Gilchrist (1994), studies across countries and identification strategies have shown that small firms respond more strongly to monetary policy, often because they have weaker balance sheets, fewer funding alternatives, and greater dependence on bank credit (Kashyap and Stein, 2000; Jiménez et al., 2012; Crouzet and Mehrotra, 2020; Ottonello and Winberry, 2020; Caglio et al., 2021; Cloyne et al., 2023). Limited outside options remain central to our interpretation because they allow incumbent banks to incorporate their own conditions into the lending relationship, as in theories of informational advantages and switching costs (Sharpe, 1990; Rajan, 1992; Ioannidou and Ongena, 2010). Our contribution is to show that the size differential remains visible after holding the borrowing firm and month fixed. Banks facing different surprises price credit differently to the same small borrower, whereas we detect no corresponding response for large firms. Cross-firm differences in borrower demand or lender substitution alone cannot generate this within-firm pattern.

The estimates identify differential loan pricing but do not distinguish among the internal bank mechanisms that produce it. Funding costs, risk limits, internal transfer prices, and balance-sheet exposures may all contribute to the response. The main result is independent of selecting one of these channels: bank-specific surprises are reflected in the relative rates that different lenders charge to the same small borrower.

6 Proprietary Equity Trading: The Surprises Have Consequences for Banks

Before turning to other pricing margins, we establish that the measured surprises carry economic weight for the banks themselves. Banks build positions ahead of policy announcements (Lucca and Moench, 2015; Bernile et al., 2016), and they enter each meeting with positions executed at different pre-event prices, which can reflect beliefs as well as hedging, inventory, client flow, and execution. If the realized bank-event transaction-price surprise captures economically meaningful exposure rather than statistical noise, banks experiencing larger absolute surprises should realize weaker trading results around the same meetings. Regulatory securities data allow us to test this prediction directly on banks' own profit and loss.

Under the Markets in Financial Instruments Regulation (MiFIR), banks report every securities transaction with the executing entity, instrument, quantity, price, and timestamp. We use proprietary transactions, executed on the banks' own account, in liquid German equities and compute for each bank and Governing Council meeting the profit over the seven calendar days up to and including

the meeting day, valuing all positions at the closing price of the meeting day.[5] Because a small number of very large trading books dominates the raw profit distribution, our main specification winsorizes the outcome, capping the most extreme profits and losses; the unadjusted raw outcome is reported alongside. The resulting panel contains 662 bank-event observations from 24 banks and 31 meetings. Because not every bank trades OIS at every event, the specification separates two margins: an indicator for observed OIS participation and, for participants, the absolute surprise intensity. Bank-events without observed proprietary positions are excluded rather than treated as zero profits. All regressions include bank and event fixed effects, so the estimates compare banks within the same meeting after absorbing the level of the aggregate market movement and permanent differences in trading scale.

Table 6 reports the result. Banks with larger absolute surprises earn significantly lower proprietary trading profits around the same meeting: the winsorized coefficient of -10.248 per percentage point implies that a 10-basis-point absolute surprise is associated with a loss of approximately EUR 1 million, with a bank-cluster WCR-S p-value of 0.016. The raw outcome yields a larger coefficient with the same sign. Participation itself carries no comparable association, so the result reflects the intensity of the surprise among active banks rather than selection into OIS trading. The surprises are thus not merely constructs from swap prices: they are mirrored in losses on the banks' own balance sheets, in a different market and a different asset class.

Table 7 decomposes the intensity into the absolute common component and the bank-specific increment, which satisfy an exact additive identity for participating banks. The role of the event fixed effects differs from Table 6 in one respect worth stating. The fixed effects absorb the common market movement itself, which affects all banks in a meeting. The common intensity in Table 7 enters interacted with OIS participation, and this interaction varies across banks within the event: it measures how much more strongly the profits of OIS-active banks co-move with the size of the common surprise than the profits of inactive banks. The common event intensity carries the association: its coefficient is -17.0 in the winsorized specification and precisely estimated, while the bank-specific increment is negative but imprecise, and the difference between the two coefficients is statistically significant. Around large policy events, OIS-active banks as a group realize losses that scale with how strongly the event surprised the market they price; the increment that separates one participant from another adds no comparable explanatory power for equity profits. The leave-own-bank-out construction reproduces this allocation (Appendix Table A.7).

The equity evidence validates the economic content of the full bank-event measure with a stake computed from banks' own transaction reports, their trading profit and loss, and it identifies which component carries that validation: common event intensity among OIS-active banks relative to

[5] We restrict attention to stocks that were constituents of the DAX index at some point during the past ten years, which ensures liquid instruments and reliable closing prices.

inactive banks. The lender-specific loan result of Section 5 neither requires nor uses this channel; it is identified from the bank-relative component within firm and month. Together, the two results begin to map where each component of the surprise matters, a pattern that the deposit market completes in the next section.

7 Corporate Deposit Rate Transmission

We next examine pass-through to overnight corporate deposit rates. Deposit pricing is central to the transmission of monetary policy into bank funding costs and the real economy (Drechsler et al., 2017; Egan et al., 2017; Drechsler et al., 2021; Wang et al., 2022), and deposit rates adjust gradually and unevenly across banks (Hannan and Berger, 1991). The wholesale deposits in our sample are observed daily and provide a considerably denser pricing margin than loan originations. They are also transacted in the unsecured money-market segment from which the euro short-term rate is computed and are therefore closely connected to the operational implementation of monetary policy. We first estimate their response to the full bank-event surprise and then separate the common and bank-relative components.

We estimate daily local projections for bank-depositor relationships. The outcome is the cumulative change from business day $t-1$ through horizon h in the rate on an active overnight deposit relationship. Bank-depositor fixed effects absorb permanent relationship heterogeneity, and depositor-by-Governing-Council-cycle fixed effects absorb common depositor conditions within the policy cycle. The specification further includes five lags. Table 8 reports the response to a 10-basis-point movement in the original bank-event surprise.

We find rapid and persistent deposit-rate pass-through. A 10-basis-point surprise increases the cumulative deposit rate by 5.7 basis points after two business days, 11.0 basis points after five days, and 12.0 basis points after ten days. The response declines gradually thereafter but remains positive through day 25. A global WCR-S max-$|t|$ test across horizons 5, 10, 15, 20, and 25 yields a p-value of 0.0097. This family test asks whether the pass-through path as a whole is zero: it compares the largest test statistic across the five horizons with its bootstrap distribution under the joint null, thereby accounting for the dependence between horizons and for multiple testing. Economically, it establishes that the deposit response is a persistent path rather than a single significant day. These estimates apply to active relationships whose rate or amount changes. Appendix diagnostics confirm that the path is robust to excluding windows that overlap with the next Governing Council event and to alternative treatments of relationship activity.

To determine which part of the surprise is associated with deposit repricing, Table 9 enters C_t and G_{bt} jointly on a fixed common sample. The coefficients are scaled to a 10-basis-point movement in the corresponding component. The decomposition compares the response to monetary policy

news shared across banks with the response to bank-relative OIS pricing revisions around the same event.

The common component is associated with a 16.8-basis-point response after five business days and remains positive through day 25. By contrast, the bank-relative coefficient is positive but imprecisely estimated at every horizon. Direct common-minus-bank-relative tests reject equality at horizons 10 through 20 and are marginally significant at day five. Thus, deposit rates respond detectably to monetary policy news shared across banks, while the estimates provide no comparable evidence for bank-relative pass-through. Because both the surprise and deposit rates originate in MMSR transactions, the decomposition also reduces the concern that the full-measure result merely reflects bank-specific pricing across the same reporting source: the detectable response appears on the common rather than the bank-relative component. The deposit data impose two restrictions worth noting. MMSR does not record depositor characteristics, so a split by depositor size analogous to the firm-size split in the loan analysis is not feasible. In addition, a same-depositor design analogous to the loan regressions is feasible only on very thin support and yields an early positive but imprecise bank-specific estimate (Appendix A.5).

The deposit results establish that the transaction-based surprise measure is economically relevant for a second bank-set price. They also reveal a sharp difference between the margins. The overnight deposits in our sample are priced off the common overnight benchmark, and it is the common component of the surprise that moves them; the bank-specific component matters most where banks price credit to firms. The deposit estimates constitute a separate transmission result rather than a first stage for the loan regressions. Across the three margins we study, a consistent division emerges: common event news moves banks' trading books and deposit rates, while the bank-specific component is priced in banks' lending to firms.

8 Event Scope: Governing Council Decisions and Speeches

Monetary policy information also arrives outside scheduled decisions. In the United States, speeches by the Federal Reserve Chair move financial markets at least as much as FOMC announcements (Swanson and Jayawickrema, 2025). For the euro area, newly assembled communication-event databases show that inter-meeting communication, including speeches by Executive Board members, has important asset-price and real effects (Istrefi et al., 2026; Altavilla et al., 2025). This literature motivates a direct scope test for our setting: does the Small-minus-Large incidence arise around bank-specific OIS price changes more generally, or is it concentrated at Governing Council decisions?

We use the ECB Executive Board and President speech events assembled in the euro area monetary policy event-study database of Altavilla et al. (2025) and measure bank-specific surprises over the same narrow three-hour windows used in that literature, so that our speech surprises

correspond to the events its authors have shown to move markets. For the speech windows we apply the leave-own-bank-out construction: with few trades per window, a bank's own post-window transactions could dominate the common benchmark, and excluding them prevents any mechanical own-bank correlation. The Governing Council surprise remains the original measure throughout; the two constructions are nearly identical for Governing Council events (Appendix Table A.2). Our main speech measure sums the available speech surprises within each bank-month; the underlying series covers 666 bank-months, 32 banks, and 49 months, with approximately two measured events per observed bank-month. We include this measure jointly with the Governing Council surprise in the strict same-borrower specification. An availability indicator and its firm-size interaction account for missing speech measurement rather than coding unobserved bank-months as zero. Since firm-month fixed effects absorb common monthly news, the comparison concerns bank-specific event components and asks whether speech-window surprises generate the same cross-lender incidence in credit pricing.

Table 10 reports the results. Panel A presents the baseline Governing Council Small-minus-Large coefficients, and Panel B includes both event classes jointly. The Governing Council coefficients remain positive and statistically significant and increase slightly to 3.095, 3.477, and 5.068 at horizons zero through two. By contrast, the speech coefficients are close to zero at horizons zero and two, negative at horizon one, and statistically insignificant throughout. Panel C directly rejects equality of the Governing Council and speech coefficients at the 5 percent level at horizons zero and one and at the 10 percent level at horizon two.

The same-borrower size differential is therefore concentrated around scheduled Governing Council decisions in our sample. This finding complements rather than contradicts evidence that speeches move asset prices and economic activity in the aggregate. Common speech news is absorbed by firm-month fixed effects, just like the common component of Governing Council surprises; our test concerns whether cross-bank speech-window revisions reproduce the lender-specific loan-pricing pattern. They do not. Bank-level speech surprises are measured over narrower windows and fewer trades than Governing Council surprises. Even so, the headline Governing Council coefficient remains strong when speech surprises enter simultaneously, and the two coefficients differ significantly at each reported horizon.

9 Conclusion

This paper constructs bank-specific monetary policy surprises from OIS transactions around ECB Governing Council meetings and studies their consequences for corporate loan pricing, banks' own equity trading, and deposit rates. The measure follows standard high-frequency identification but retains the identity of the transacting bank. Aggregating the transactions reproduces established

euro area surprise series, while the bank-event panel reveals substantial variation among banks exposed to the same policy decision. Firm-month fixed effects absorb the common event component but leave this bank-relative variation, permitting a same-borrower comparison of the total response of each borrower group, which a single aggregate shock cannot deliver.

The loan estimates reveal a pronounced difference by firm size. Banks facing larger surprises reprice credit more strongly for small firms than for large firms borrowing in the same month. Moving from the 10th to the 90th percentile of the within-event surprise distribution implies 22 to 38 basis points more repricing for small firms in the first and second complete months after the decision, while the response for large firms remains close to zero. The differential emerges in the first complete month after the decision, is stable after controlling for loan amount, instrument type, and maturity, and is reproduced by the leave-own-bank-out measure. Positive surprises raise small-firm rates relatively more and negative surprises lower them relatively more. Because the regressions compare different lenders to the same firm in the same month, cross-firm differences in credit demand or access to lenders cannot explain this pattern.

The surprises are economically consequential beyond the loan market. Banks with larger absolute surprises experience losses in their proprietary equity trading around the same meetings, of approximately EUR 1 million per 10 basis points, and active overnight deposit rates respond rapidly and persistently to the full surprise measure. In both margins, the detectable response is concentrated in common event news rather than the bank-relative component. Moreover, speech-window surprises do not reproduce the Governing Council Small-minus-Large coefficient when both event types enter jointly. Taken together, the results show that common monetary policy news moves banks' trading books and deposit rates, while granular surprise variation is especially informative for the pricing of credit to firms around scheduled policy decisions.

For the analysis of monetary policy transmission, the findings shift attention from borrower characteristics to the lending relationship itself. It is well established that small firms react more strongly to monetary policy, and this pattern is usually attributed to their weaker balance sheets and outside options. Our results show that an important part of the transmission to small firms is decided at the level of the individual bank: after the same decision, the same small firm pays measurably different rates depending on which lender it faces and on how that lender was surprised. Small firms therefore bear not only stronger average transmission but also the dispersion of their lenders' assessments, a burden that large firms escape because competition ties their rates to a common benchmark. In the euro area, where small and medium-sized enterprises dominate the firm population and rely primarily on bank credit, this lender-driven dispersion is a dimension of transmission that a common surprise measure cannot reveal, even in loan-level data. The regulatory data needed to measure it already exist within the Eurosystem.

The measurement results carry a complementary message. The concentration of OIS activity

among a limited number of institutions means that aggregate surprise measures summarize pricing revisions that can differ materially across banks. The granular measure does not replace aggregate shocks; it reveals the cross-bank variation that aggregation leaves unobserved and thereby extends the set of questions that high-frequency identification can answer.

Our estimates describe one euro area policy cycle and a concentrated set of OIS-active banks, and the strict same-borrower estimates rely on 10 or 11 bank clusters, which is why we report few-cluster and horizon-family inference together with extensive stability tests. The measurement approach itself carries over to any setting in which transaction repositories identify participants and execution times. What it delivers is a simple reorientation: aggregate surprises measure what a policy decision did to the market, while granular surprises show how the banks that price policy were positioned when the decision arrived, and what that meant for the firms borrowing from them. On both counts, the answer of this paper is that monetary policy transmission is not uniform: banks facing different surprises reprice loans differentially, and it is small firms that carry this incidence.

References

Altavilla, C., L. Brugnolini, R. Gürkaynak, R. Motto, and G. Ragusa (2019). Measuring euro area monetary policy. *Journal of Monetary Economics 108*(C), 162–179.

Altavilla, C., F. Canova, and M. Ciccarelli (2020). Mending the broken link: Heterogeneous bank lending rates and monetary policy pass-through. *Journal of Monetary Economics 110*, 81–98.

Altavilla, C., R. S. Gürkaynak, T. Kind, and L. Laeven (2025). Monetary transmission with frequent policy events. Working Paper Series 3157, European Central Bank.

Altavilla, C., R. S. Gürkaynak, and R. Quaedvlieg (2024). Macro and micro of external finance premium and monetary policy transmission. *Journal of Monetary Economics 147*, 103634.

Amiti, M. and D. E. Weinstein (2018). How much do idiosyncratic bank shocks affect investment? evidence from matched bank-firm loan data. *Journal of Political Economy 126*(2), 525–587.

Andrade, P., G. Gaballo, E. Mengus, and B. Mojon (2019). Forward guidance and heterogeneous beliefs. *American Economic Journal: Macroeconomics 11*(3), 1–29.

Bahaj, S., R. Czech, S. Ding, and R. Reis (2025). The market for inflation risk. Discussion Paper 20157, Centre for Economic Policy Research.

Bauer, M. and E. Swanson (2023). An alternative explanation for the "Fed information effect". *American Economic Review 113*(3), 664–700.

Bauer, M. D., C. E. Pflueger, and A. Sunderam (2024). Perceptions about monetary policy. *Quarterly Journal of Economics 139*(4), 2227–2278.

Begenau, J., M. Piazzesi, and M. Schneider (2015). Banks' risk exposures. Working Paper 21334, National Bureau of Economic Research. Revised August 2025.

Bernile, G., J. Hu, and Y. Tang (2016). Can information be locked up? informed trading ahead of macro-news announcements. *Journal of Financial Economics 121*(3), 496–520.

Bordalo, P., N. Gennaioli, Y. Ma, and A. Shleifer (2020). Overreaction in macroeconomic expectations. *American Economic Review 110*(9), 2748–2782.

Caglio, C. R., R. M. Darst, and Ş. Kalemli-Özcan (2021). Collateral heterogeneity and monetary policy transmission: Evidence from loans to SMEs and large firms. Working Paper 28685, National Bureau of Economic Research.

Cloyne, J., C. Ferreira, M. Froemel, and P. Surico (2023). Monetary policy, corporate finance, and investment. *Journal of the European Economic Association 21*(6), 2586–2634.

Coibion, O. and Y. Gorodnichenko (2012). What can survey forecasts tell us about information rigidities? *Journal of Political Economy 120*(1), 116–159.

Crouzet, N. and N. R. Mehrotra (2020). Small and large firms over the business cycle. *American Economic Review 110*(11), 3549–3601.

D'Acunto, F., J. Gao, L. Liu, K. Lu, Z. Wang, and J. Yang (2025). Subjective expectations and financial intermediation. Discussion Paper 20280, Centre for Economic Policy Research.

de Bondt, G. (2005). Interest rate pass-through: Empirical results for the euro area. *German Economic Review 6*(1), 37–78.

Degryse, H., O. De Jonghe, S. Jakovljević, K. Mulier, and G. Schepens (2019). Identifying credit supply shocks with bank-firm data: Methods and applications. *Journal of Financial Intermediation 40*, 100813.

Degryse, H. and S. Ongena (2005). Distance, lending relationships, and competition. *Journal of Finance 60*(1), 231–266.

Drechsler, I., A. Savov, and P. Schnabl (2017, 05). The deposits channel of monetary policy. *The Quarterly Journal of Economics 132*(4), 1819–1876.

Drechsler, I., A. Savov, and P. Schnabl (2021). Banking on deposits: Maturity transformation without interest rate risk. *Journal of Finance 76*(3), 1091–1143.

ECB (2025, April). Euro money market study 2024. Technical report, European Central Bank, Frankfurt am Main.

Egan, M., A. Hortaçsu, and G. Matvos (2017). Deposit competition and financial fragility: Evidence from the US banking sector. *American Economic Review 107*(1), 169–216.

Fecht, F., B. Imbierowicz, and S. Greppmair (2025). Corporate deposit reallocation and bank funding stability. Working paper, Deutsche Bundesbank.

Gabaix, X. (2011). The granular origins of aggregate fluctuations. *Econometrica 79*(3), 733–772.

Gabaix, X. and R. S. J. Koijen (2024). Granular instrumental variables. *Journal of Political Economy 132*(7), 2274–2303.

Gertler, M. and S. Gilchrist (1994). Monetary policy, business cycles, and the behavior of small manufacturing firms. *The Quarterly Journal of Economics 109*(2), 309–340.

Gertler, M. and P. Karadi (2015). Monetary policy surprises, credit costs, and economic activity. *American Economic Journal: Macroeconomics 7*(1), 44–76.

Gomez, M., A. Landier, D. Sraer, and D. Thesmar (2021). Banks' exposure to interest rate risk and the transmission of monetary policy. *Journal of Monetary Economics 117*, 543–570.

Gürkaynak, R., B. Sack, and E. Swanson (2005). Do actions speak louder than words? the response of asset prices to monetary policy actions and statements. *International Journal of Central Banking 1*(1), 55–93.

Hannan, T. H. and A. N. Berger (1991). The rigidity of prices: Evidence from the banking industry. *American Economic Review 81*(4), 938–945.

Heider, F., F. Saidi, and G. Schepens (2019). Life below zero: Bank lending under negative policy rates. *Review of Financial Studies 32*(10), 3728–3761.

Hoffmann, P., S. Langfield, F. Pierobon, and G. Vuillemey (2019). Who bears interest rate risk? *Review of Financial Studies 32*(8), 2921–2954.

Ioannidou, V. and S. Ongena (2010). "time for a change": Loan conditions and bank behavior when firms switch banks. *Journal of Finance 65*(5), 1847–1877.

Istrefi, K., F. Odendahl, and G. Sestieri (2026). ECB communication and its impact on financial markets. *European Economic Review 188*, 105410.

Jarociński, M. (2022). Central bank information effects and transatlantic spillovers. *Journal of International Economics 139*, 103683.

Jarociński, M. and P. Karadi (2020, April). Deconstructing monetary policy surprises: The role of information shocks. *American Economic Journal: Macroeconomics 12*(2), 1–43.

Jeenas, P. (2023). Firm balance sheet liquidity, monetary policy shocks, and investment dynamics. Working Paper 1409, Barcelona School of Economics.

Jiménez, G., S. Ongena, J.-L. Peydró, and J. Saurina (2012). Credit supply and monetary policy: Identifying the bank balance-sheet channel with loan applications. *American Economic Review 102*(5), 2301–2326.

Jiménez, G., S. Ongena, J.-L. Peydró, and J. Saurina (2014). Hazardous times for monetary policy: What do twenty-three million bank loans say about the effects of monetary policy on credit risk-taking? *Econometrica 82*(2), 463–505.

Jordà, Ò. (2005). Estimation and inference of impulse responses by local projections. *The American Economic Review 95*(1), 161–182.

Kashyap, A. K. and J. C. Stein (2000). What do a million observations on banks say about the transmission of monetary policy? *American Economic Review 90*(3), 407–428.

Khwaja, A. I. and A. Mian (2008). Tracing the impact of bank liquidity shocks: Evidence from an emerging market. *American Economic Review 98*(4), 1413–1442.

Kuttner, K. N. (2001). Monetary policy surprises and interest rates: Evidence from the Fed funds futures market. *Journal of Monetary Economics 47*(3), 523–544.

Lucca, D. O. and E. Moench (2015). The pre-FOMC announcement drift. *Journal of Finance 70*(1), 329–371.

MacKinnon, J. G., M. Ø. Nielsen, and M. D. Webb (2023). Fast and reliable jackknife and bootstrap methods for cluster-robust inference. *Journal of Applied Econometrics 38*(5), 671–694.

Mankiw, N. G., R. Reis, and J. Wolfers (2004). Disagreement about inflation expectations. *NBER Macroeconomics Annual 18*, 209–248.

Milgrom, P. and N. Stokey (1982). Information, trade and common knowledge. *Journal of Economic Theory 26*(1), 17–27.

Miranda-Agrippino, S. and G. Ricco (2021). The transmission of monetary policy shocks. *American Economic Journal: Macroeconomics 13*(3), 74–107.

Nakamura, E. and J. Steinsson (2018). High-frequency identification of monetary non-neutrality: The information effect. *Quarterly Journal of Economics 133*(3), 1283–1330.

Ottonello, P. and T. Winberry (2020). Financial heterogeneity and the investment channel of monetary policy. *Econometrica 88*(6), 2473–2502.

Paravisini, D., V. Rappoport, and P. Schnabl (2023). Specialization in bank lending: Evidence from exporting firms. *Journal of Finance 78*(4), 2049–2085.

Patton, A. J. and A. Timmermann (2010). Why do forecasters disagree? lessons from the term structure of cross-sectional dispersion. *Journal of Monetary Economics 57*(7), 803–820.

Petersen, M. A. and R. G. Rajan (1995). The effect of credit market competition on lending relationships. *Quarterly Journal of Economics 110*(2), 407–443.

Rajan, R. G. (1992). Insiders and outsiders: The choice between informed and arm's-length debt. *Journal of Finance 47*(4), 1367–1400.

Romer, C. D. and D. H. Romer (2000). Federal reserve information and the behavior of interest rates. *American Economic Review 90*(3), 429–457.

Romer, C. D. and D. H. Romer (2004). A new measure of monetary shocks: Derivation and implications. *The American Economic Review 94*(4), 1055–1084.

Sastry, K. A. (2026). Disagreement about monetary policy. *American Economic Journal: Macroeconomics 18*(2), 486–522.

Schmeling, M., A. Schrimpf, and S. A. M. Steffensen (2022). Monetary policy expectation errors. *Journal of Financial Economics 146*(3), 841–858.

Sharpe, S. A. (1990). Asymmetric information, bank lending, and implicit contracts: A stylized model of customer relationships. *Journal of Finance 45*(4), 1069–1087.

Swanson, E. T. (2021). Measuring the effects of federal reserve forward guidance and asset purchases on financial markets. *Journal of Monetary Economics 118*, 32–53.

Swanson, E. T. and V. Jayawickrema (2025). Speeches by the Fed chair are more important than FOMC announcements: An improved high-frequency measure of U.S. monetary policy shocks. Working paper, University of California, Irvine.

Vilerts, K., S. Anyfantaki, K. Beņkovskis, S. Bredl, M. Giovannini, F. M. Horky, V. Kunzmann, T. Lalinský, A. Lampousis, E. Lukmanova, F. Petroulakis, and K. Zutis (2025). Details matter: Loan pricing and transmission of monetary policy in the euro area. Working Paper 3078, European Central Bank.

Wang, Y., T. M. Whited, Y. Wu, and K. Xiao (2022). Bank market power and monetary policy transmission: Evidence from a structural estimation. *Journal of Finance 77*(4), 2093–2141.

Figures

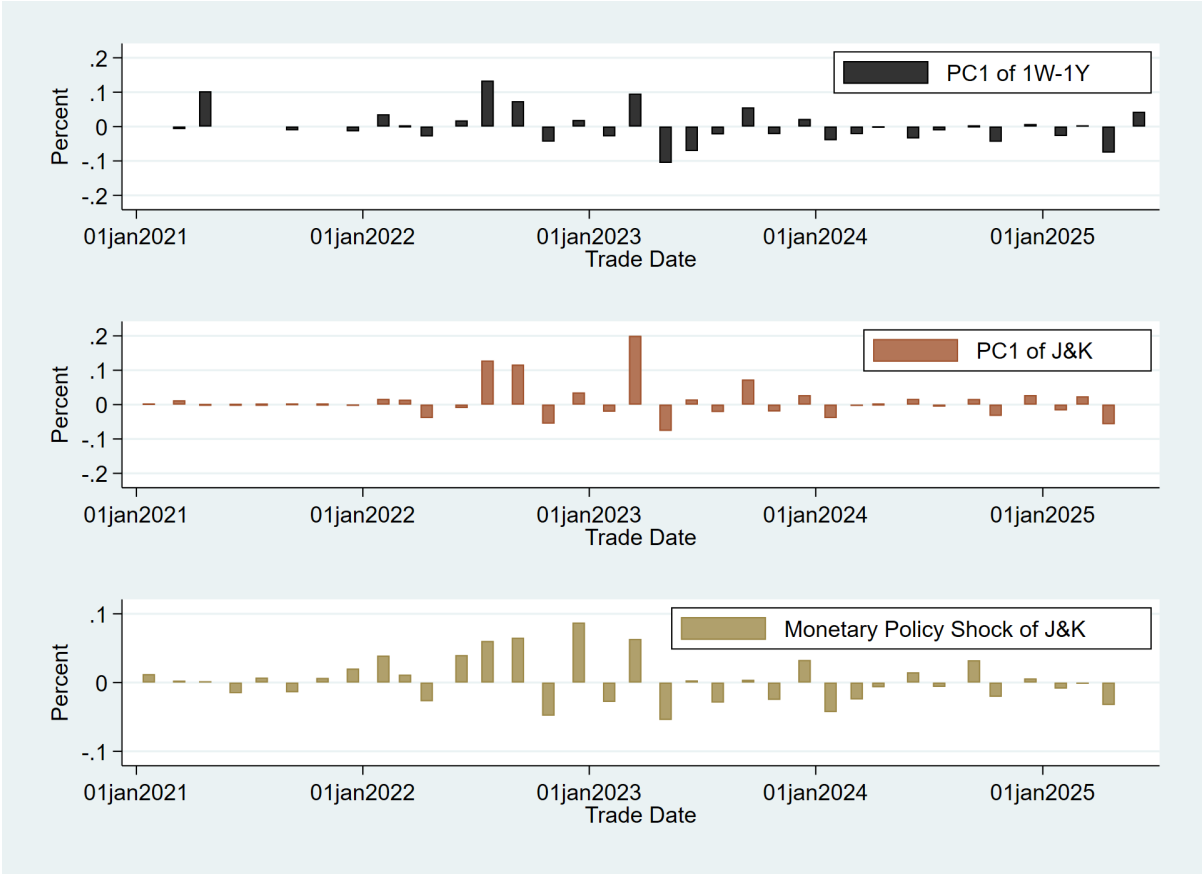
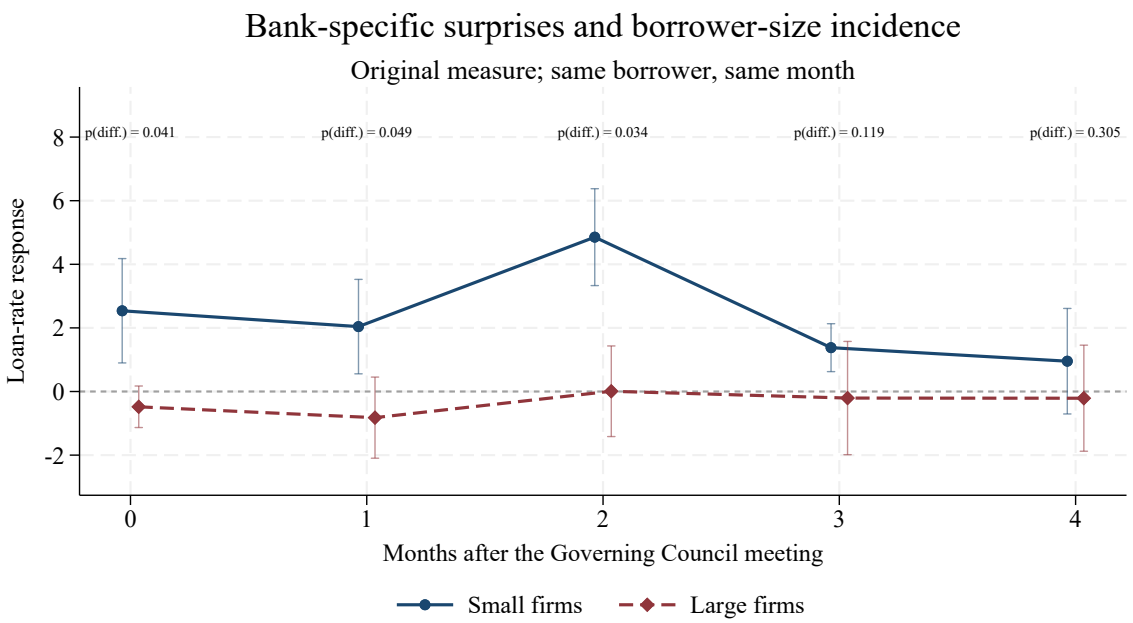


Figure 1. Aggregate Monetary Policy News from Bank-Level Transactions

Notes: The figure compares the first principal component of transaction-based OIS rate changes, aggregated across banks before extraction, with the Jarociński-Karadi interest-rate surprise and their pure monetary policy shock. All series are expressed in percentage points. Summary statistics of the bank-event measure, including its support and within-event dispersion, are reported in Table 1.

(a) Total surprise slopes by firm size



Whiskers are 95% bank-clustered confidence intervals. p(diff.) tests Small minus Large using WCR-S. Joint horizon-family p = 0.031.

(b) Implied differential repricing of small firms, 90th versus 10th percentile surprise

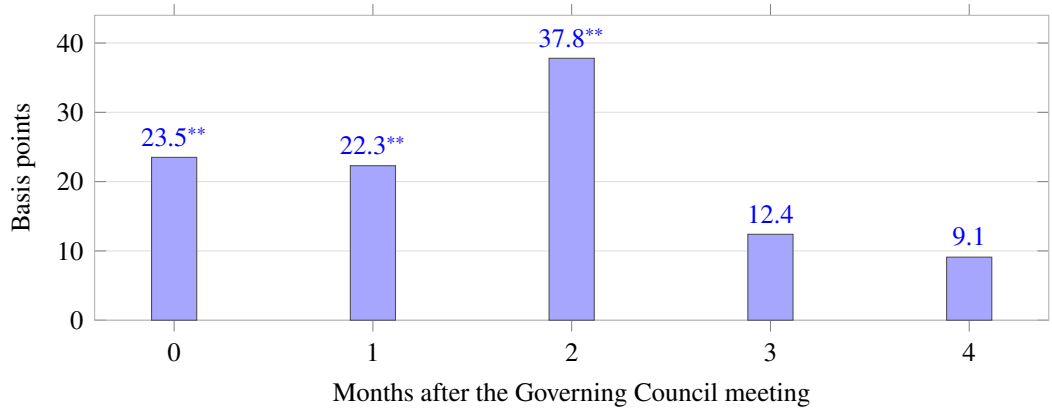


Figure 2. Bank-Specific Surprises and Borrower-Size Incidence

Notes: Panel A: The dependent variable is the cumulative loan-rate change from month $t-1$ through horizon h. Points report the total surprise slopes for small and large firms from one joint same-borrower regression at each horizon (Table 2, Panel B); whiskers are 95 percent conventional bank-clustered confidence intervals. The displayed p-values test equality of the two slopes using bank-cluster WCR-S with 9,999 Webb draws; the synchronized horizon-family p-value is 0.031. Panel B: Each bar scales the Small-minus-Large contrast at that horizon by the horizon-specific within-event interdecile (90th minus 10th percentile) range of the surprise. It shows the additional cumulative repricing of a small firm whose lender experienced a 90th-percentile surprise relative to the same firm borrowing from a 10th-percentile-surprise lender in the same month. Stars carry over from the direct WCR-S equality tests: ** $p < 0.05$.

Tables

Table 1. External Benchmark, Scale, and Support of Granular Monetary Policy Surprises

Panel A. Agreement with established aggregate surprise measures

Construction from regulatory OIS transactions	JK rate surprise	JK pure monetary policy surprise
Transactions aggregated before principal-component extraction	0.80	0.72
Completed bank-event surprise series aggregated	0.73	0.69

Panel B. Original measure

Variable	N	Mean	SD	Min	Max
Transaction-specific surprise (bp)	4,127	-0.91	11.06	-62.15	90.67
Bank-event surprise (bp)	366	0.00	6.10	-34.23	20.68
Absolute bank-event surprise (bp)	366	4.44	4.18	0.01	34.23

Panel C. Bank-event support

Support statistic	Original measure
Observed bank-events	366
Banks	31
Governing Council events with observed bank-events	32
Average within-event standard deviation (bp)	3.79

Notes: Panel A reports correlations. The first row aggregates transactions before extracting the principal component and is the closest comparison with conventional aggregate high-frequency identification. The second row aggregates the completed bank-event series. JK denotes Jarociński and Karadi (2020). Panels B and C describe the original measure; the event calendar contains 36 meetings, and bank-event surprises can be constructed for 32 of them. The original measure compares each bank's pre-announcement OIS transactions with a common volume-weighted post-announcement benchmark matched by event, maturity band, and trade direction.

Table 2. Loan Repricing: Conventional Saturation versus the Same-Borrower Design

	Monthly horizon				
	(0)	(1)	(2)	(3)	(4)
Panel A. Bank × Firm fixed effects only					
A.1 Pooled model: one joint regression per horizon					
Common component	0.056	0.378	1.298	2.253	2.230
	[0.952]	[0.768]	[0.451]	[0.344]	[0.402]
Bank-specific component	-0.153	-0.171	0.207	0.326	0.163
	[0.753]	[0.780]	[0.835]	[0.811]	[0.919]
A.2 Firm-size interaction model: one joint regression per horizon					
Common component × Small firm	0.004	0.826	1.088	-1.171	-1.005
	[0.996]	[0.413]	[0.442]	[0.522]	[0.392]
Bank-specific component × Small firm	0.258	-1.117	0.395	-6.276	-4.525
	[0.827]	[0.120]	[0.899]	[0.176]	[0.197]
Panel B. Same-borrower design: Bank × Firm and Firm × Month fixed effects					
Surprise × Small firm	2.539**	2.041	4.853**	1.379	0.953
	[0.044]	[0.102]	[0.033]	[0.121]	[0.184]
Surprise × Large firm	-0.479	-0.821	0.009	-0.206	-0.210
	[0.154]	[0.156]	[0.989]	[0.884]	[0.834]
Post-estimation Small-minus-Large contrast from Panel B					
Difference in pass-through	**3.018**	**2.863**	**4.845**	1.585	1.164
p-value for H_0: Small = Large	**0.041**	**0.049**	**0.034**	0.119	0.305
Implied P90-P10 differential (bp)	23.5	22.3	37.8	12.4	9.1
Observations	26,129	19,985	18,712	17,975	16,854
Banks	11	11	10	11	11
Bank × Firm FE	Yes	Yes	Yes	Yes	Yes
Firm × Month FE (Panel B only)	Yes	Yes	Yes	Yes	Yes
Full lags	Yes	Yes	Yes	Yes	Yes

Notes: The dependent variable is the cumulative change in the amount-weighted loan rate from month $t-1$ through horizon h. Both panels use the same horizon-specific samples. Panel A omits firm-month fixed effects, so the common and bank-specific surprise components can be estimated jointly; A.1 reports the pooled slopes and A.2 their Small-firm interaction coefficients, each relative to the omitted Large-firm group, from a separate regression. Panel B adds firm-month fixed effects, which absorb the common component and all borrower conditions shared by a firm's lenders in a month; each column is one joint regression estimating separate total slopes for mutually exclusive lagged firm-size groups. Large firms have lagged annual turnover above the sample median; all other observed firms are small. The contrast block is a post-estimation linear test, not an additional regressor. Square brackets report restricted wild-cluster-bootstrap (WCR-S) p-values based on 9,999 Webb draws, clustered by bank for bank-specific coefficients and by event for common-component coefficients. The P90-P10 row scales the direct contrast by the horizon-specific within-event interdecile range of the surprise. Horizon zero uses all loans originated in the event month and is therefore not a strictly post-decision estimate; Table 3 isolates the post-decision part of the event month. Stars refer to the Panel B estimates and denote * $p < 0.10$, ** $p < 0.05$, and *** $p < 0.01$.

Table 3. Timing of Differential Loan Repricing

	Rest of event month		Through first full month	Through second full month
	Strict	Inclusive	$h = 1$	$h = 2$
	(1)	(2)	(3)	(4)
Small-minus-Large contrast	-0.098	0.310	2.863**	4.845**
	[0.901]	[0.821]	[0.049]	[0.034]
Observations	10,292	11,073	19,985	18,712
Banks	11	11	11	10
Bank × Firm FE	Yes	Yes	Yes	Yes
Firm × Month FE	Yes	Yes	Yes	Yes
Full lags	Yes	Yes	Yes	Yes

Notes: Each column is a separate Khwaja-Mian regression; the reported estimand is the post-estimation difference between the Small- and Large-firm surprise slopes. The event month is month t. Strict compares the amount-weighted rate on loans originated strictly after the Governing Council date in month t with month $t - 1$; Inclusive also includes loans originated on the Governing Council date. Through first full month is the rate in $t + 1$ minus the rate in $t - 1$; Through second full month is the rate in $t + 2$ minus the rate in $t - 1$. Columns (1) and (2) use a different rate aggregate and a thinner support than column (0) of Table 2, which averages over all loans of the event month. Square brackets report bank-cluster WCR-S p-values based on 9,999 Webb draws.

Table 4. Differential Loan Repricing with Contract-Level Controls

Control set, estimated separately	Monthly horizon				
	(0)	(1)	(2)	(3)	(4)
Baseline: Small minus Large	3.018**	2.863**	4.845**	1.585	1.164
	[0.041]	[0.049]	[0.034]	[0.119]	[0.305]
Plus amount, instrument type, mean maturity	2.997**	2.801**	4.922**	1.531	1.173
	[0.042]	[0.045]	[0.038]	[0.122]	[0.307]
Plus amount, instrument type, maturity bins	2.962**	2.891**	4.873**	1.439	1.132
	[0.039]	[0.049]	[0.035]	[0.120]	[0.305]
Observations, richest controls	26,125	19,985	18,708	17,971	16,854
Banks	11	11	10	11	11
Bank × Firm FE	Yes	Yes	Yes	Yes	Yes
Firm × Month FE	Yes	Yes	Yes	Yes	Yes
Full lags	Yes	Yes	Yes	Yes	Yes

Notes: Each row and horizon is a separate Khwaja-Mian regression. Every entry is the post-estimation Small-minus-Large contrast; rows differ only in the listed contract controls. Mean maturity is linear residual maturity; maturity bins replace that term with maturity-category indicators. The observations row reports the richest-control sample. Square brackets report bank-cluster WCR-S p-values based on 9,999 Webb draws.

Table 5. Borrower-Size Incidence by the Sign of the Bank-Specific Surprise

	Monthly horizon		
	(0)	(1)	(2)
Panel A. Coefficients from one joint sign model per horizon			
Positive Surprise × Small firm, relative to Large	4.358*	1.671**	4.760**
	[0.083]	[0.046]	[0.040]
Negative Surprise × Small firm, relative to Large	1.085	4.612**	3.443*
	[0.467]	[0.047]	[0.071]
Observations	26,129	19,985	18,712
Banks	11	11	10
Positive bank-event cells	41	41	34
Negative bank-event cells	45	42	40
Bank × Firm FE	Yes	Yes	Yes
Firm × Month FE	Yes	Yes	Yes
Full lags	Yes	Yes	Yes
Panel B. Post-estimation equality test			
p-value for H_0: Positive slope = Negative slope	0.195	0.164	0.180

Notes: Each column is one joint regression including the positive and negative parts of the signed original bank-event surprise and their Small-firm interactions. Reported coefficients are Small-minus-Large slope differences. Positive and Negative refer to the sign of the bank-specific transaction-price surprise, not to policy regimes. Panel B tests equality of the two slope differences within the same regression. Square brackets and Panel-B entries report bank-cluster WCR-S p-values based on 9,999 Webb draws.

Table 6. Bank-Event Surprise Intensity and Proprietary Equity Trading Profits

Dependent variable: 7-day proprietary equity profit (EUR mn)	Winsorized (1)	Raw (2)
Absolute surprise intensity × OIS participation	−10.248**	−16.126*
	(5.021)	(9.048)
	[0.0159]	[0.0522]
OIS participation	0.398	0.465
	[0.232]	[0.383]
Observations	662	662
Banks	24	24
Events	31	31
Bank FE	Yes	Yes
Event FE	Yes	Yes
Two-margin design	Yes	Yes

Notes: Each column is one regression of banks' 7-day proprietary equity trading profit around a Governing Council meeting on the absolute bank-event surprise interacted with OIS participation. The surprise is expressed in percentage points; in Column 1, a 10-basis-point absolute surprise implies a loss of approximately EUR 1.025 million. Parentheses report bank-clustered standard errors for the surprise-intensity coefficient. Square brackets report bank-cluster WCR-S p-values based on 9,999 Webb draws for that coefficient and conventional bank-clustered p-values for OIS participation. The two-margin design distinguishes the intensive surprise exposure of observed OIS participants from participation itself; bank-events without observed proprietary positions are excluded rather than treated as zero profits. Stars denote * $p < 0.10$, ** $p < 0.05$, and *** $p < 0.01$.

Table 7. Common Event Intensity, Bank-Specific Increments, and Equity Trading Profits

Dependent variable: 7-day proprietary equity profit (EUR mn)	Winsorized (1)	Raw (2)
Absolute common component × OIS participation	−17.002***	−30.602**
	(6.316)	(12.654)
	[0.0048]	[0.0152]
Bank-specific increment to absolute intensity	−6.727	−8.581
	(4.838)	(7.023)
	[0.1825]	[0.2443]
Observations	662	662
Banks	24	24
Events	31	31
Bank FE	Yes	Yes
Event FE	Yes	Yes
OIS participation included	Yes	Yes
Exact additive intensity decomposition	Yes	Yes
Post-estimation cross-coefficient and joint tests		
Common minus bank-specific increment	−10.275*	−22.021**
	[0.0665]	[0.0147]
Joint p-value: both components equal zero	0.0202	0.0270

Notes: Each column is one joint regression. The two components enter together and satisfy the exact identity $|S_{bt}| = |C_t| + [|S_{bt}| - |C_t|]$ for OIS-participating bank-event cells. The second component is the bank-specific increment to absolute surprise intensity; it can be negative when bank-specific pricing offsets common event news and is not $|S_{bt} - C_t|$. Parentheses report bank-clustered standard errors and square brackets report bank-cluster WCR-S p-values based on 9,999 Webb draws. Event fixed effects absorb the common outcome level, so the common coefficient identifies how OIS participants' relative profits vary with common event magnitude. The final two rows are post-estimation tests rather than additional coefficient estimates. Stars denote * $p < 0.10$, ** $p < 0.05$, and *** $p < 0.01$.

Table 8. Monetary Policy News and Active Overnight Deposit Rates

	Business-day horizon					
	(2)	(5)	(10)	(15)	(20)	(25)
Cumulative rate response to a 10-bp original bank-event surprise						
Original bank-event Surprise	5.720**	11.015***	12.036***	10.373***	7.718***	6.763***
	[0.0248]	[0.0001]	[0.0000]	[0.0000]	[0.0001]	[0.0000]
Observations	39,309	37,998	36,351	35,694	35,007	33,909
Banks	24	23	24	24	24	23
Events	31	31	31	30	30	30
Bank × Depositor FE	Yes	Yes	Yes	Yes	Yes	Yes
Depositor × GC-cycle FE	Yes	Yes	Yes	Yes	Yes	Yes
Full lags	Yes	Yes	Yes	Yes	Yes	Yes
Post-estimation family test						
Global horizon-family p-value			0.0097			

Notes: Each column is a separate local-projection regression. The dependent variable is the cumulative change from business day $t-1$ through horizon h in the interest rate of an actively repriced unsecured overnight deposit relationship. Coefficients are scaled to a 10-basis-point original bank-event surprise. Square brackets report pointwise conventional bank-clustered p-values. The global horizon-family p-value is from a bank-cluster WCR-S max-$|t|$ test, based on 9,999 synchronized Webb draws and a common fixed-support sample, of the joint null that coefficients at horizons 5, 10, 15, 20, and 25 are zero.

Table 9. Common Event News and Bank-Relative Deposit-Rate Repricing

	Business-day horizon				
	(5)	(10)	(15)	(20)	(25)
Common event component	16.769***	16.871***	14.846***	13.011***	9.023***
	[0.0021]	[0.0027]	[0.0016]	[0.0025]	[0.0043]
Bank-relative component	2.182	1.871	1.129	1.304	1.545
	[0.7691]	[0.8077]	[0.8631]	[0.7053]	[0.8695]
Observations	22,028	22,028	22,028	22,028	22,028
Banks	22	22	22	22	22
Events	29	29	29	29	29
Bank × Depositor FE	Yes	Yes	Yes	Yes	Yes
Depositor × GC-cycle FE	Yes	Yes	Yes	Yes	Yes
Full lags	Yes	Yes	Yes	Yes	Yes
Post-estimation cross-coefficient equality test					
Common minus bank-relative	14.587*	14.999**	13.717**	11.706***	7.478
	[0.0712]	[0.0440]	[0.0101]	[0.0007]	[0.1598]

Notes: Each column is one joint regression in which the common and bank-relative surprise components enter together on a fixed common sample. Coefficients are scaled to a 10-basis-point movement in the relevant component. Square brackets report bank-cluster WCR-S p-values based on 9,999 Webb draws. The final block is the post-estimation difference between the two coefficients and tests H_0: Common = Bank-relative; it is not a third estimated coefficient.

Table 10. Governing Council and Speech Surprises in Same-Borrower Loan Repricing

	Monthly horizon		
	Event month	First full month	Two-month cumulative
	$h = 0$	$h = 1$	$h = 2$
Panel A. Baseline model: one regression per horizon			
GC Surprise × Small firm, relative to Large	3.018**	2.863**	4.845**
	[0.041]	[0.049]	[0.034]
Panel B. Joint GC-Speech model: one regression per horizon			
GC Surprise × Small firm, relative to Large	3.095**	3.477**	5.068**
	[0.0318]	[0.0387]	[0.0307]
Speech Surprise × Small firm, relative to Large	0.168	-1.074	0.116
	[0.7793]	[0.3314]	[0.9244]
Observations	26,129	19,985	18,712
Banks	11	11	10
Observations with measured Speech Surprise	18,532	14,062	13,109
Bank × Firm FE	Yes	Yes	Yes
Firm × Month FE	Yes	Yes	Yes
Full lags and Speech-availability controls	Yes	Yes	Yes
Panel C. Post-estimation tests from the joint model			
GC minus Speech coefficient	2.927**	4.551**	4.951*
	[0.0265]	[0.0365]	[0.0504]
Joint p-value: GC and Speech both zero	0.1070	0.0775	0.0378

Notes: Panel A reports the Governing Council (GC) Small-minus-Large contrast from the baseline model. Panel B is a separate joint regression in which granular GC surprises and monthly sums of bank-specific surprises measured over narrow three-hour speech windows, together with their Small-firm interactions, enter simultaneously. The GC surprise is the original measure; the speech surprises use a leave-own-bank-out post-window benchmark because few trades are observed per speech window. Missing speech measurement is represented by an availability indicator and its Small-firm interaction rather than an observed zero. Panel C reports pointwise tests at each horizon, first for equality of the GC and Speech contrasts and then for their joint significance. Square brackets report bank-cluster WCR-S p-values based on 9,999 Webb draws. Firm-by-Month Fixed Effects absorb common monthly news, so the table compares bank-specific event components rather than aggregate speech effects.

Appendix

Granular Monetary Policy Surprises

A.1 Data Construction and Sample Formation

A.1.1 OIS Cleaning and Matching

The OIS sample is constructed in six steps. First, we extract new and valid OIS transactions reported in MMSR on the 36 Governing Council days and remove novations. Second, we retain centrally cleared contracts. Third, unique transaction identifiers link MMSR to EMIR records that supply execution timestamps and both reported counterparties. Fourth, we deduplicate mirror reports and reconcile trade dates, counterparties, and directions; when minor rate discrepancies remain, MMSR economic terms take precedence. Fifth, we flag material inconsistencies in rates, notionals, execution dates, start dates, or maturities and apply conservative filters. Sixth, we retain spot and near-spot forward OIS contracts with maturities from one week to ten years, classify pay- and receive-fixed directions, and assign trades to pre-event, event, and post-event windows using the release and press-conference schedule.

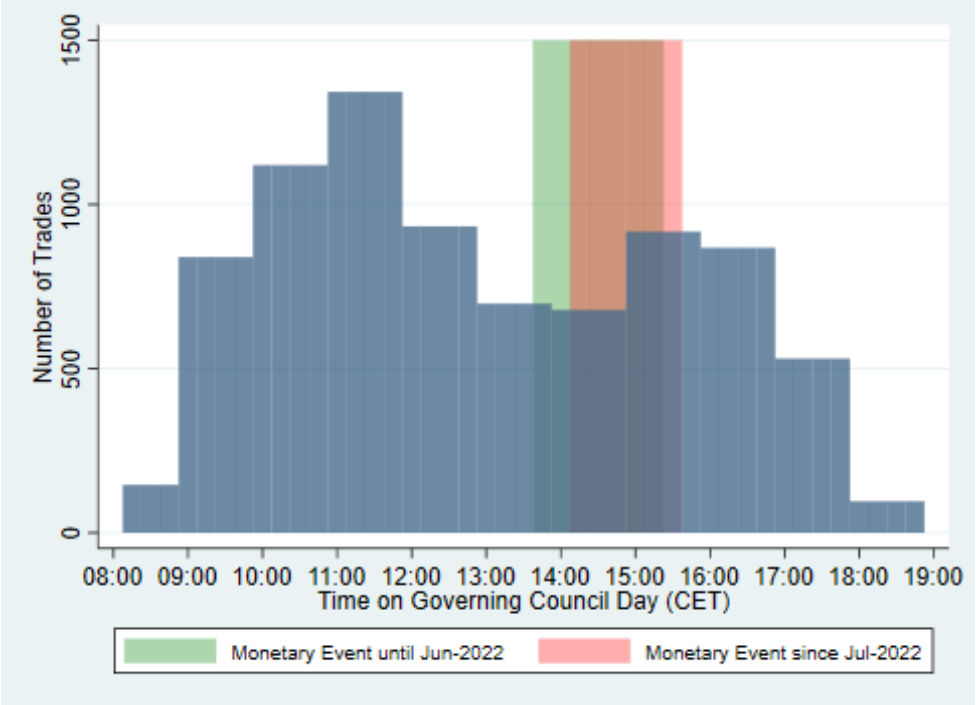
The linked panel retains approximately 91 percent of the relevant MMSR transactions for which an EMIR timestamp match is sought. The construction uses no unobserved bank-event as a zero surprise. A transaction-specific surprise requires both a pre-event transaction and a post-event benchmark in the same event, narrow maturity band, and direction. This requirement explains the decline from 5,210 pre-event trades to 4,127 transaction-specific surprises. Aggregating those observations produces 366 measured bank-events from 31 banks at 32 of the 36 meetings.

Figure A.1 documents why the event windows use the available trading day rather than a narrow quote window. Activity is greatest earlier in the day, declines near the release, and resumes after the press conference. Figure A.2 shows event-day volume and rate dispersion by maturity and policy phase.

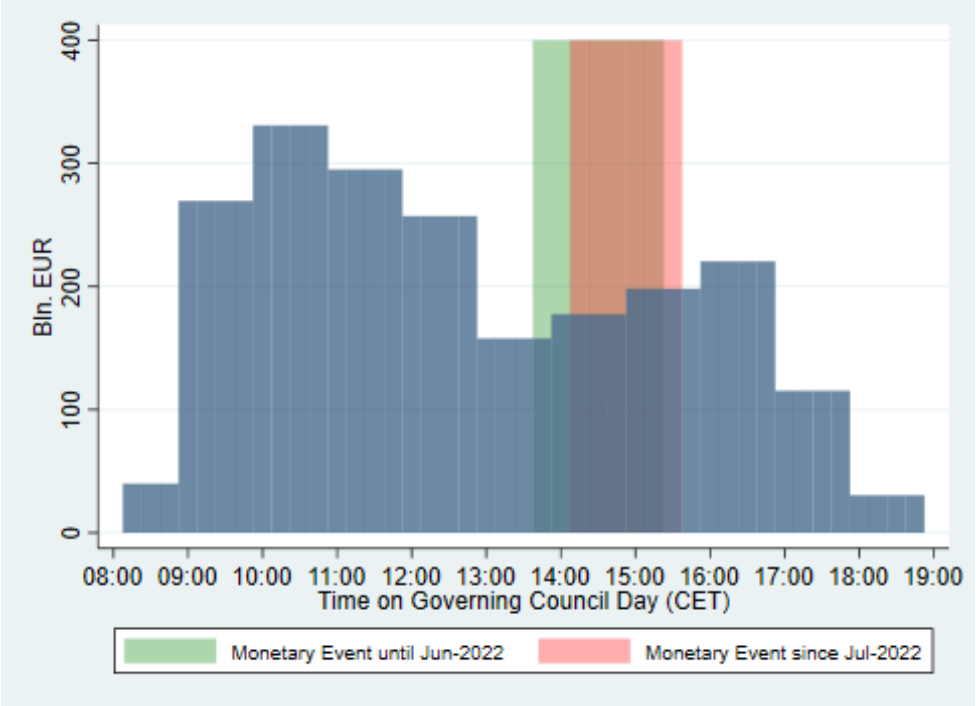
Table A.1. Sample Formation and Effective Estimation Support

Stage	Observations	Banks	Events
Cleaned OIS transactions on Governing Council days	8,171	31	36
Pre-event OIS transactions	5,210	31	36
Post-event OIS transactions	1,726	31	36
Transaction-specific surprises	4,127	31	36
Completed bank-event surprises	366	31	32
Predetermined-characteristics regression (Table A.3)	359	29	29
Pooled same-borrower loan design, horizon 0	28,834	12	–
Strict loan design, horizon 0	26,129	11	–
Strict loan design, horizon 2	18,712	10	–
Active deposit design, horizon 2 business days	39,309	24	31
Fixed-sample deposit component decomposition	22,028	22	29
Equity two-margin design	662	24	31

Notes: Observation units change by stage: transactions, transaction-specific surprises, bank-events, bank-firm-months, and bank-depositor-days. A dash indicates that an event count is not defined for the corresponding observation unit. The completed bank-event measure spans 32 of the 36 meetings in the event calendar. The table distinguishes broad source coverage from the effective bank clusters in the strict same-borrower design.



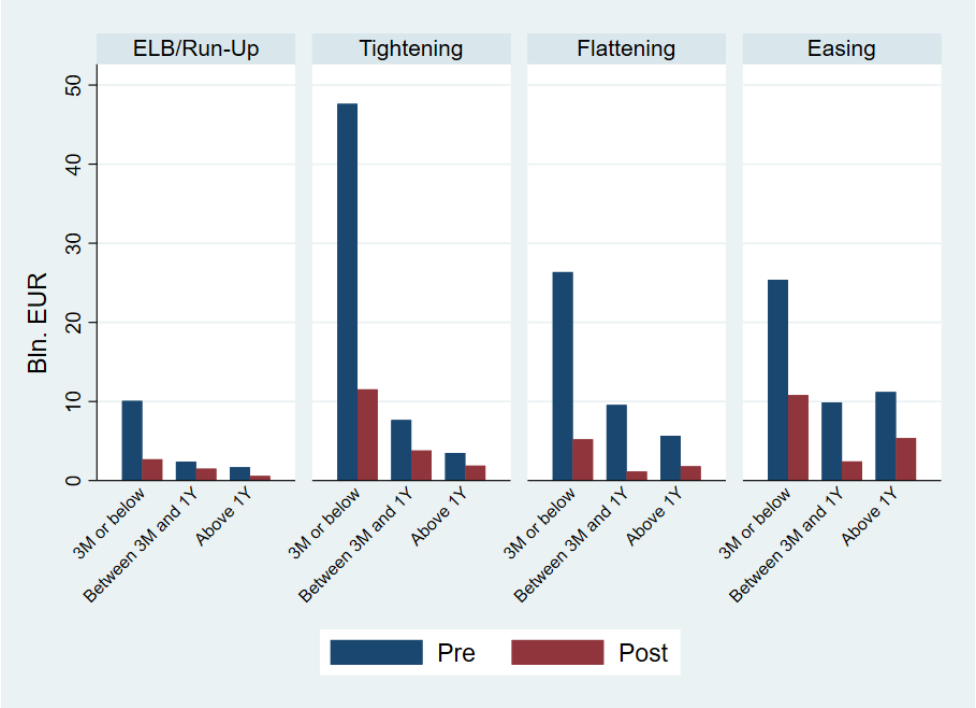
(a) Number of transactions



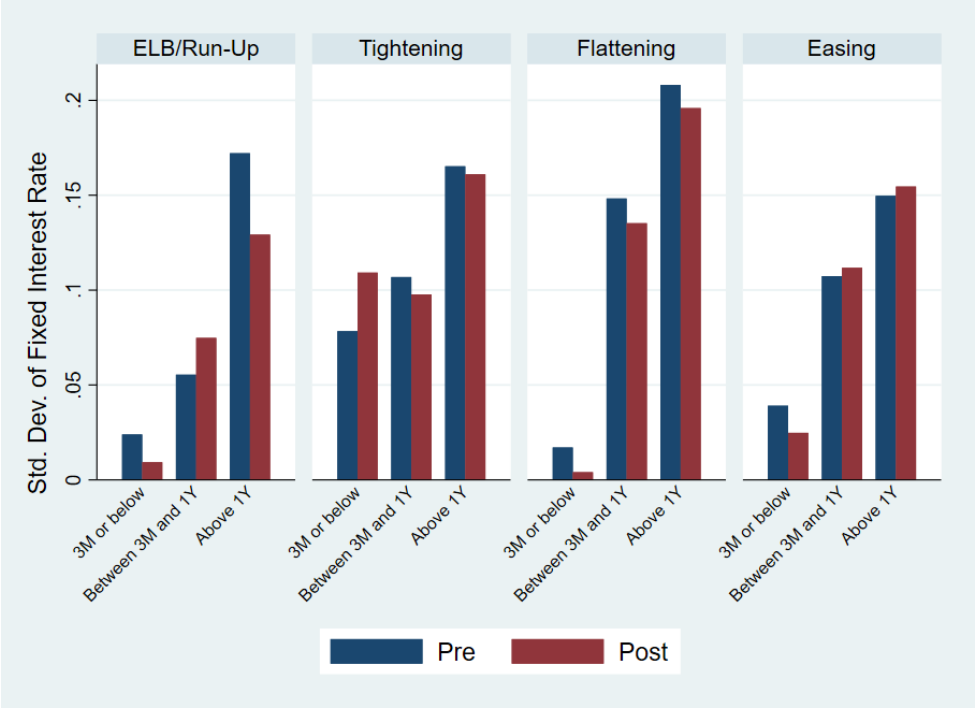
(b) Transaction volume

Figure A.1. Intraday OIS Activity on Governing Council Days

Notes: The panels show event-day OIS transaction counts and total notional volume by time of day. Event-window colors reflect the change in the ECB release schedule in July 2022. Sample: January 2021 to June 2025.



(a) Volume by maturity and policy phase



(b) Rate dispersion by maturity and policy phase

Figure A.2. OIS Trading Support across Maturities and Policy Phases

Notes: Panel A reports pre- and post-event volume. Panel B reports the volume-weighted standard deviation of fixed OIS rates. Maturity buckets are mutually exclusive. Policy phases follow the dates indicated in the underlying event file.

A.1.2 Loan and Deposit Samples

AnaCredit observations are limited to newly originated euro-denominated fixed-rate term loans and non-revolving credit lines with maturity no longer than one year. We require positive agreed rates, coincident inception and settlement dates, nondefaulted debtors, and consistent lender and borrower identifiers. Within a bank-firm-month, rates are weighted by instrument amount. Contract-control specifications retain amount, instrument type, and exact or binned residual maturity before aggregation. The strict Khwaja-Mian sample then requires at least two lenders to the same firm-month, observed group assignment, the relevant lagged rates, and identifying variation after fixed effects.

The MMSR deposit sample is restricted to unsecured overnight transactions in which the reporting bank borrows from a nonfinancial corporation. Same-day settlement, positive volume, fixed rate, and absence of notice or option features are required. Ultimate-parent identifiers consolidate counterparties. Multiple bank-firm-day reports are volume weighted. The active sample records a pricing observation only when rate or volume differs from the previous business day. No-overlap robustness truncates the outcome before the next Governing Council meeting.

A.2 Measurement Robustness

A.2.1 Leave-Own-Bank-Out Construction

The original measure uses a common post-event benchmark to avoid treating post-announcement portfolio adjustment as part of the surprise. A bank could nevertheless contribute to that common benchmark. The LOBO construction removes bank b's own post-announcement transactions before calculating the matched benchmark for its pre-event trades. This changes only the benchmark; all outcome definitions and local-projection specifications remain fixed.

Table A.2. Original and Leave-Own-Bank-Out Construction on Common Support

Panel A. Support

	Original	LOBO
Observed bank-events	366	361
Banks	31	29
Governing Council events	32	31
Within-event standard deviation (bp)	3.79	4.29

Panel B. Agreement on exact common support

Agreement statistic	Estimate
Correlation in levels	0.963
Within-event correlation	0.916
Rank correlation	0.982
Same within-event sign	93.4%
Mean absolute within-event difference	0.69 bp
Original-on-LOBO slope	1.036
	(0.050)

Notes: Leave-own-bank-out (LOBO) excludes bank b's own post-announcement transactions from the common post-event benchmark used to construct its surprise. The final row reports the regression coefficient and standard error recovered from the reported 95 percent confidence interval.

The level correlation is 0.963, the within-event correlation is 0.916, and 93.4 percent of common observations retain their within-event sign. The average absolute within-event difference is 0.69 basis points. These statistics show that each bank's direct contribution to the common post-event price is quantitatively small.

A.2.2 Predetermined Bank Characteristics

Appendix Table A.3 regresses signed and absolute bank-event surprises on lagged bank characteristics. With event fixed effects, the full-vector joint p-values are 0.268 for signed surprises and 0.265 for absolute surprises; wholesale deposit dependence is individually significant at the 10 percent level in the signed specification, and log assets is negatively associated with absolute surprise intensity in the event-fixed-effect specification; both are individual associations within a jointly insignificant vector. Adding bank fixed effects isolates within-bank variation and yields joint p-values of 0.731 and 0.864. The results indicate limited joint predictive power of the observed bank characteristics both across banks within an event and over time within banks. They are evidence against systematic sorting on these observables rather than evidence of random assignment.

Table A.3. Predetermined Bank Characteristics and Granular Monetary Policy Surprises

Panel A. Signed surprise with event fixed effects

	Baseline (1)	Full (2)
L.Log(Assets), standardized	-0.127	-0.213
	[0.394]	[0.203]
L.Capital ratio, standardized	-0.002	-0.123
	[0.988]	[0.488]
L.Wholesale deposit dependence, standardized		-0.371*
		[0.083]
L.NFC deposit dependence, standardized		0.136
		[0.381]
Joint WCR-S p-value	0.694	0.268
Observations	359	359
Banks / Events	29 / 29	29 / 29
Event FE	Yes	Yes

Panel B. Absolute surprise with event fixed effects

	Full
L.Log(Assets), standardized	-0.437**
	[0.043]
L.Capital ratio, standardized	-0.375
	[0.170]
L.Wholesale deposit dependence, standardized	-0.253
	[0.435]
L.NFC deposit dependence, standardized	-0.077
	[0.707]
Joint WCR-S p-value	0.265
Observations	359
Banks / Events	29 / 29
Event FE	Yes

Panel C. Joint predictive content with bank and event fixed effects

Dependent variable	Absorbed fixed effects	Joint WCR-S p-value
Signed bank-event Surprise	Bank and Event	0.731
Absolute bank-event Surprise	Bank and Event	0.864

Notes: The sample is the common support of the original and leave-own-bank-out measures (361 bank-events); event fixed effects drop two singleton events, leaving 359. All bank characteristics are lagged and standardized, so coefficients are basis points per one-standard-deviation change. Square brackets report bank-cluster WCR-S p-values based on 9,999 Webb draws; joint tests refer to the full vector of characteristics. Stars denote * $p < 0.10$ and ** $p < 0.05$.

A.3 Loan Robustness and Inference

A.3.1 LOBO Replication of the Headline Estimand

Appendix Table A.4 estimates the LOBO measure on the exact original loan outcome samples. The early Small-minus-Large contrasts are 2.580, 1.995, and 3.969, with WCR-S p-values of 0.043, 0.046, and 0.036. The conclusion therefore does not rely on a bank entering its own post-event benchmark.

Table A.4. Leave-Own-Bank-Out Surprise and Differential Loan Repricing

	Monthly horizon				
	(0)	(1)	(2)	(3)	(4)
LOBO Surprise × Small firm	2.086**	1.405	4.037**	1.772*	0.886*
	[0.047]	[0.102]	[0.034]	[0.083]	[0.061]
LOBO Surprise × Large firm	-0.494	-0.591	0.068	-0.082	-0.082
	[0.114]	[0.140]	[0.804]	[0.892]	[0.912]
Observations	26,129	19,985	18,712	17,975	16,854
Banks	11	11	10	11	11
Bank × Firm FE	Yes	Yes	Yes	Yes	Yes
Firm × Month FE	Yes	Yes	Yes	Yes	Yes
Full lags	Yes	Yes	Yes	Yes	Yes
Post-estimation Small-minus-Large contrast					
Difference in pass-through	**2.580**	**1.995**	**3.969**	1.855	0.968
p-value for H_0: Small = Large	**0.043**	**0.046**	**0.036**	0.100	0.249
Implied P90-P10 differential (bp)	21.6	16.7	33.3	15.6	8.1

Notes: LOBO excludes each bank's own post-announcement transactions from its common benchmark and is estimated on the exact original outcome sample used in Table 2. Square brackets and equality-test entries report bank-cluster WCR-S p-values based on 9,999 Webb draws.

A.3.2 Pooled Same-Borrower Estimates

Appendix Table A.5 separates the pooled average from borrower-size incidence. Firm-fixed-effect estimates become positive at later horizons, whereas the same-borrower coefficients remain near zero and are sometimes negative. The difference between both specifications is positive at every horizon but is statistically insignificant both pointwise and over the horizon family. At horizon one, the confidence intervals and minimum detectable effects also show that the pooled same-borrower estimate is not an equivalence result. The table therefore documents the identifying difference between the two specifications without implying that conventional estimates significantly overstate average pass-through.

Table A.5. Pooled Loan-Rate Pass-Through under Alternative Fixed Effects

	Monthly horizon				
	(0)	(1)	(2)	(3)	(4)
Panel A. Original measure					
Firm-FE coefficient	-0.064	0.126	0.890	1.664	1.600
Strict KM coefficient	-0.432	-0.802	0.080	-0.255	-0.250
Firm-FE minus KM	0.368	0.928	0.810	1.918	1.850
	[0.474]	[0.405]	[0.515]	[0.390]	[0.378]
Family p-value for Firm-FE minus KM			0.391		
Panel B. LOBO measure					
Firm-FE coefficient	-0.045	0.313	1.084	1.926	1.943
Strict KM coefficient	-0.452	-0.594	0.121	-0.110	-0.104
Firm-FE minus KM	0.406	0.906	0.963	2.036	2.047
	[0.412]	[0.183]	[0.392]	[0.358]	[0.362]
Family p-value for Firm-FE minus KM			0.397		
Panel C. Informativeness of the strict KM estimate at horizon 1, per 10-bp surprise					
Original: estimate / 95% CI / 80% MDE (bp)			−8.0 / [−21.0, 5.0] / 18.0		
LOBO: estimate / 95% CI / 80% MDE (bp)			−5.9 / [−12.8, 0.9] / 9.5		

Notes: Firm-FE models retain bank-firm fixed effects but omit firm-month fixed effects. Strict KM models include both. Square brackets report bank-cluster WCR-S p-values for the direct difference between the two model coefficients on harmonized support. MDE denotes the two-sided minimum detectable effect at 80 percent power. The table addresses average pass-through; it does not estimate the Small-minus-Large incidence in the main text.

A.3.3 Few-Cluster and Family-Wise Inference

Table A.6 summarizes the inference procedures for the loan result, each of which answers one concrete question. The pointwise WCR-S tests address the small number of bank clusters: with 10 or 11 banks, conventional clustered standard errors reject far too often, and the restricted bootstrap addresses this by imposing the null when resampling. The family tests address multiple testing across horizons: because the result is a path over five horizons, they ask whether all five Small-minus-Large differences are jointly zero. The common-bank version holds the set of banks fixed across horizons and yields p-values of 0.031 for the original measure and 0.032 for LOBO; the conservative fixed-row-support version additionally holds the observations fixed across horizons and yields 0.082 and 0.081. The omission rows address the concern that a single institution or a single meeting drives the result: the early differential remains positive when any one bank or any one event is removed.

Table A.6. Few-Cluster and Family-Wise Inference for the Loan Incidence Result

	Original	LOBO
Pointwise WCR-S p-values for Small minus Large		
$h = 0$	0.041	0.043
$h = 1$	0.049	0.046
$h = 2$	0.034	0.036
Horizon-family tests, $h = 0, \ldots, 4$ *jointly zero*		
Synchronized common-bank family	0.031	0.032
Conservative fixed-row-support family	0.082	0.081
Stability under omission, $h = 0, 1, 2$		
Positive under every single-bank omission	Yes	Yes
Positive under every single-event omission	Yes	Yes

Notes: WCR-S denotes the restricted wild cluster bootstrap with 9,999 Webb draws, clustered by bank. The pointwise rows report the p-values of the direct Small-minus-Large equality tests at each horizon. The family rows test the joint null that all five horizon contrasts are zero: the synchronized common-bank family uses horizon-specific observation samples and a common set of banks, while the conservative fixed-row-support family additionally fixes the observation support across horizons (10,674 observations and eight banks at every horizon). The omission rows re-estimate the early contrasts excluding one bank or one Governing Council event at a time.

A.3.4 Continuous Size and Additional Support

Replacing the median split with standardized lagged log turnover produces a negative size gradient at every horizon: repricing becomes weaker as firms become larger. The common horizon-family p-values are 0.272 for the original measure and 0.235 for LOBO. The continuous specification therefore supports the direction of the main result but is less precisely estimated. We retain the binary split as the primary specification because it provides an interpretable direct contrast between mutually exclusive firm-size groups.

A.4 Equity Robustness

The proprietary-trading outcome is constructed from MiFIR securities-transaction reports. For each bank and Governing Council meeting, we cumulate the signed profit on own-account equity transactions over the seven calendar days up to and including the meeting day, valuing positions at the closing price of the meeting day. The sample covers liquid German equities, and the outcome is winsorized in the main specification with the raw outcome reported alongside. Structural zeros are handled through the two-margin design in the main text: OIS participation enters as its own margin, and absolute surprise intensity is measured only for participating bank-events. Of the 475 bank-events in the equity panel without a measured surprise, 428 involve no qualifying pre-announcement OIS trade and enter through the participation margin; the remaining 47, for which a pre-event trade exists but no compatible post-event benchmark, are excluded, which reduces the panel from 709 to 662

bank-events.

Appendix Table A.7 re-estimates the intensity decomposition with the leave-own-bank-out surprise. The allocation is unchanged: the absolute common component remains negative and precisely estimated, the bank-specific increment remains imprecise, and the difference between the two coefficients is statistically significant in both the winsorized and the raw specification. The equity result therefore does not depend on a bank contributing to its own post-announcement benchmark.

Table A.7. Leave-Own-Bank-Out Robustness of the Equity Intensity Decomposition

Dependent variable: 7-day proprietary equity profit (EUR mn)	Winsorized (1)	Raw (2)
Absolute common component × OIS participation	−15.719***	−32.105**
	(5.665)	(13.941)
	[0.0025]	[0.0105]
Bank-specific increment to absolute intensity	−2.701	2.610
	(2.389)	(3.998)
	[0.2430]	[0.5523]
Observations	657	657
Banks	24	24
Events	31	31
Bank FE	Yes	Yes
Event FE	Yes	Yes
OIS participation included	Yes	Yes
Exact additive intensity decomposition	Yes	Yes
Post-estimation cross-coefficient and joint tests		
Common minus bank-specific increment	−13.018**	−34.715**
	[0.0363]	[0.0190]
Joint significance of both components	[0.0122]	[0.0165]

Notes: The leave-own-bank-out construction removes each bank from the post-announcement benchmark used to construct its surprise. The two components enter jointly and satisfy the same exact intensity identity as in Table 7. Parentheses report bank-clustered standard errors and square brackets report bank-cluster WCR-S p-values based on 9,999 Webb draws. The final blocks are post-estimation tests. The five-observation difference from Table 7 reflects leave-own-bank-out support rather than a changed outcome definition. Stars denote * $p < 0.10$, ** $p < 0.05$, and *** $p < 0.01$.

A.5 Deposit Robustness

Additional measurement and timing checks support the deposit result. The original and LOBO dynamic paths have a correlation of 0.9994, and the LOBO horizon-family p-value is 0.0055. Truncating responses before the next Governing Council event leaves the positive path intact. Reconstructing the active-repricing filter on 371,225 relationship-date keys yields no robust evidence

that the surprise predicts selection into a changed rate or amount, and raw and winsorized outcomes are nearly identical. Extending the outcome to all carried-forward deposit reports, by contrast, produces smaller and less precise responses. The evidence therefore applies to active overnight deposit repricing rather than every outstanding balance.

The strict bank-relative same-depositor-day design has considerably thinner support: 1,167 observations, 16 banks, and 15 treated depositor-days. It produces an early positive estimate but is sensitive to event and horizon-family inference. We therefore base the deposit interpretation on the joint common and bank-relative decomposition in Table 9, which uses the denser relationship panel.

A.6 Event Scope

The speech exercise uses narrow three-hour windows and includes a measurement-availability margin. The underlying monthly speech series contains 666 observed bank-months from 32 banks and 49 months, with approximately two speech events per observed bank-month. A joint Governing Council and speech specification for deposit rates has insufficient high-dimensional support: only 62 observations remain at horizon one and 47 at horizon two. We therefore focus on the loan comparison in Table 10, which directly tests whether bank-relative speech surprises reproduce the same-borrower incidence found around Governing Council decisions. This comparison is distinct from the aggregate effects of speeches on asset prices or real activity.